

This file is part of the following work:

**Trevathan, Jarrod (2007) *Privacy and security in online auctions*. PhD Thesis,
James Cook University.**

Access to this file is available from:

<https://doi.org/10.25903/rnvf%2Daz70>

Copyright © 2007 Jarrod Trevathan

The author has certified to JCU that they have made a reasonable effort to gain permission and acknowledge the owners of any third party copyright material included in this document. If you believe that this is not the case, please email

researchonline@jcu.edu.au

Privacy and Security in Online Auctions

Thesis submitted by
Jarrod Trevathan

for the Degree of Doctor of Philosophy
in the School of Mathematics, Physics and Information Technology at
James Cook University

supervisor: Associate Professor Wayne Read

jarrod.trevathan@jcu.edu.au
wayne.read@jcu.edu.au

March 2007

Declaration

I declare that this thesis is my own work and has not been submitted in any form for another degree or diploma at any university or other institute of tertiary education. Information derived from the published and unpublished work of others has been acknowledged in the text, and a list of references is given.

Jarrod Trevathan
March, 2007

Statement on Access to this Thesis

I, the under-signed, the author of this work, understand that James Cook University will make this work available for use within the University Library, and via the Australian Digital Thesis Network, for use elsewhere.

I understand that as an unpublished work, a thesis has significant protection under the Copyright Act. I do not wish to place any restriction on access to this thesis. However, any use of its content must be acknowledged and could potentially be restricted by future patents.

Jarrod Trevathan
March, 2007

Acknowledgments

Firstly, I would like to thank Wayne Read for his supervision throughout this thesis.

Next, I would like to thank my parents Ian and Bev Trevathan, and brother Shane, for their support. Thank you Sharith for helping me edit my publications.

Meow to Mr Gab.

Muchas gracias para mi querida amiga Andrea.

Accolade is also due to those who directly or indirectly contributed this thesis.

Abstract

Buying and selling online is inherently insecure. Misuse of an individual's personal information is now the leading concern among those who engage in e-commerce. This thesis examines privacy and security issues in online auctions. Various auction fraud issues are investigated, and several novel counter measures proposed. An online auction server was constructed to aid in developing these security measures. This allowed investigation and testing in a controlled environment. The research results include:

1. A complete model for conducting secure and anonymous online auctions;
2. A method for detecting a fraudulent bidding practice referred to as shill bidding;
3. Autonomous bidding agents which bid maliciously. (Used to test the ability of the proposed security mechanisms.);
4. A complete model for conducting secure and anonymous online share trading; and
5. Several alternate proposals for auction clearing algorithms.

The proposed security mechanisms have been implemented on the online auction server. Results are given as simulated and practical tests. In addition, the auction server's software design is documented. Many of the techniques discussed in this thesis can be readily applied to commercial online auctions.

Contents

1	Introduction	1
1.1	Aims	2
1.2	Methodology	2
1.3	Results	3
1.4	Thesis Organisation	3
2	Online Auction Software 1	5
2.1	RAS: A System for Supporting Research in Online Auctions	5
3	Auction Privacy and Security Issues	15
3.1	Security, Anonymity and Trust in Electronic Auctions	15
3.2	Design Issues for Electronic Auctions	24
4	Privacy and Security in English Auctions	33
4.1	Secure Online English Auctions	33
5	Fraudulent Auctioning Practices	45
5.1	Undesirable and Fraudulent Behaviour in Online Auctions	45
5.2	Detecting Shill Bidding in Online English Auctions	56
5.3	Detecting Collusive Shill Bidding	71
6	Software Bidding Agent Security	82
6.1	A Simple Shill Bidding Agent	82
6.2	A New Approach to Avoiding the Local Extrema Trap	90
6.3	An Adaptive Shill Bidding Agent	97
7	Privacy and Security in Continuous Double Auctions	106
7.1	Privacy and Security Concerns for Online Share Trading	106
7.2	An Anonymous and Secure Continuous Double Auction Scheme	114
8	Online Auction Software 2	125
8.1	A Software Architecture for Continuous Double Auctions	125
8.2	Variable Quantity Market Clearing Algorithms	138
9	Conclusions	146

Chapter 1

Introduction

Online Auctions are extremely popular. Once the domain of highly skilled negotiators, online auctions have made auctioning accessible to everyone, regardless of who they are. A novice can buy everyday items such as groceries and clothes, or can compete for rarities and collectables such as rock memorabilia and antiques. Likewise, a seller has access to virtually a worldwide consumer base. It also seems that non-conventional items can be sold via online auctions, which would probably not be sold anywhere else! For example, auctions for the following items have been held:

1. Mold on a sandwich that resembles the Virgin Mary;
2. A man willing to bang his head on a door until he is unconscious; and
3. A French Fry shaped like the Nike logo.

However, despite the overwhelming benefits and hype, there is a sinister and dark reality to auctioning online. Auction fraud is one of the fastest growing forms of Internet-based crime. Participants are anonymous and can engage in undesirable and fraudulent behaviour in an attempt to gain an unfair advantage. For example, the seller may misrepresent or not deliver an item. Likewise, a bidder can refuse to pay, or have his/her bid forged. Furthermore, the Auctioneer could block bids or influence the auction in a manner that maximises its revenue. Various cryptographic solutions have been proposed to fix many of these problems. However, most of these schemes are not suited to the auction style that is most commonly used online.

Certain types of bidding behaviour can also be used to influence the auction in an undesirable manner. One such type of behaviour is *shilling*, where the seller introduces fake bids into the auction in order to inflate the price. Shilling is prohibited in online auctions, however, it still continues to occur. Solutions to detect and/or prevent shill bidding tend to be outside the realm of cryptography. While online auctioneers claim to monitor their auctions for bad behaviour, there are no published methods on how to detect shill bidding. The problem is compounded with the advent of *software bidding agents*.

Software bidding agents are used to bid on a human's behalf. The Artificial Intelligence community has suggested that agent-based negotiation could eventually replace all human input. However, such a claim is flawed, as agents can also act in a fraudulent manner. While much research has been conducted into improving the performance of bidding agents, little attention has been given to the security implications. A bidding agent can be designed maliciously so that it harms the auction in some of the previously mentioned ways. This is a serious concern now that bidding agents are used for trading shares online.

Online share trading is a popular offshoot of more conventional online auctions. An individual can submit a buy or sell order to a broker, who then enters it into the share market. The privacy and security issues in online auctions are also manifest in online share trading. Although tightly regulated, there is even less security than online auctions. The extra parties involved in the auctioning process (i.e., broker, share market, regulatory authority, etc.) further complicate the privacy and security requirements. Limited attention has been paid to these issues, or the special auction type employed by this application.

There are many auction types including Vickrey, Dutch, Japanese, Combinatorial, etc. Existing auction security literature has mainly concentrated on Vickrey or sealed bid auctions. However, the most popular online auction type employed online resembles an *English* auction (e.g., eBay). In an English auction, one seller offers an item to several bidders where the highest price wins. In contrast, online share trading uses an auction type referred to as a *Continuous Double Auction* (CDA). A CDA has many buyers and sellers continually trading a commodity. The method for matching buy and sell bids (referred to as *clearing*) is much more sophisticated than English auctions. The privacy and security requirements for English auctions and CDAs differ dramatically from previously studied auction types. This thesis concentrates solely on English auctions and CDAs.

1.1 Aims

This thesis investigates privacy, security and fraud issues in online English and Continuous Double auctions. The main objective is to understand the characteristics of fraudulent auction behaviour, and propose mechanisms to combat it. With regard to the aforementioned problems, the research goals include:

1. Provide a security model for auctioning online that protects a bidder's personal information;
2. Develop methods to detect certain types of auction fraud (i.e., shilling);
3. Explore the security implications of agent based negotiation;
4. Construct a model for securely and anonymously trading shares online; and
5. Devise and evaluate alternate software mechanisms for clearing CDAs.

1.2 Methodology

Many of the issues investigated in the research are **illegal** to perform for real in commercial online auctions. This makes it difficult to assess the extent of fraudulent activity, and understand its nature. Furthermore, there is no way to evaluate the effectiveness of newly proposed security mechanisms. To accomplish the research goals, an online auction server has been constructed. This allows for experimentation in a controlled environment, without the risk of prosecution.

The server was used to test security theories and gauge practical performance and efficiency. The server can perform both simulated and real auctions. It has a software bidding agent interface, which allows auctions to be fully automated (i.e., no human input required). The auction server's software components are documented throughout the thesis. Descriptions are given regarding how tests were conducted, and their results.

This thesis does not condone the use of any of the discussed themes outside the scope of the research.

1.3 Results

With regard to the aims stated in section 1.1, the thesis results are as follows:

1) Online auction privacy and security issues have been extensively investigated. From this, a complete online auction security model for English auctions has been devised. This model uses cryptographic mechanisms to ensure bid authenticity, and provides verification that everyone has followed the auction protocol correctly. An individual remains anonymous provided they don't repudiate having made a bid. In the event of bid repudiation, two independent parties can work together to trace the bid's owner.

2) A novel method to detect shill bidding in online English auctions is presented. Bidders are issued with a score based on their bidding behaviour. The score indicates the likelihood that the bidder is engaging in shilling. This can be used by other bidders to decide whether they want to participate in auctions held by a particular seller. The scheme is resilient in that it can detect colluding bidders and sellers that attempt to thwart the system.

3) Two malicious bidding agents that engage in shilling behaviour have been constructed. A *simple shill bidding agent* inserts fake bids to inflate the price until it becomes too risky to continue. A *adaptive shill agent* uses knowledge from a series of auctions with substitutable items to revise its strategy. This is achieved using novel prediction methods. The agents allow us to understand a fraudulent bidder's nature, and help refine the proposed shill detection techniques.

4) Privacy and security issues involved in online share trading are investigated, and a comprehensive set of security requirements are given for this auction type (i.e., CDAs). The online English auction security model is extended to encompass CDAs and application specific details pertinent to online share trading. This is the first secure auction scheme to specifically address share markets.

5) Several new CDA market clearing algorithms are proposed, which are more efficient than the existing methods used in financial markets. These algorithms employ techniques such as waiting, subsidisation and prioritising to achieve a higher trade volume. Efficiency is measured in terms of the number of bids and amount of quantity matched. Comparisons are drawn between the optimal offline algorithm and the proposed online clearing algorithms.

As previously mentioned, the performance of the devised techniques has been tested using the online auction server. Test results are given and scrutinised in terms of existing literature. Many of the results can be practically applied in commercial online auctions.

1.4 Thesis Organisation

This thesis is by publication. Each chapter contains independent sections with published or publishable papers. The collection of papers forms a chapter's theme. Relevant literature is reviewed at the start of each paper (an extensive literature review of auction security by the thesis author, can be found in Trevathan 2004 [1]). This thesis draws from e-Commerce, Information Security, Software Engineering and Artificial Intelligence.

The organisation of the thesis is as follows. First up, the software platform for conducting online English auctions is introduced in Chapter 2. Chapter 3 gives an overview of the security and anonymity problems with auctioning online, and is the main thrust behind the security considerations of the research.

Chapter 4 presents an online English auction security model that seeks to remedy the identified privacy and security concerns. Chapter 5 discusses fraudulent bidding behaviour and proposes methods to detect shill bidding. Chapter 6 explores the security implications for agent-based negotiation, and presents a bidding agent that bids in a fraudulent manner. Chapter 7 discusses online share trading, and its similarities to online auctions in terms of privacy and security. An anonymous and secure CDA scheme for trading shares online is presented. Chapter 8 examines the software components required for conducting CDAs, and contrasts the performance of several different online market clearing algorithm proposals. Chapter 9 provides some concluding remarks and avenues for future work.

This thesis is designed to be read sequentially from start to finish, as later topics build upon previously introduced concepts. However, after reading Chapters 2 and 3, the reader can skip ahead if desired. English auctions are addressed in Chapters 4, 5 and 6. CDAs are covered in Chapters 7 and 8.

Chapter 2

Online Auction Software 1

An online auction server was constructed to aid in testing the research proposals. This chapter presents an introductory paper describing the auction server's software characteristics. This paper forms the basis for later chapters on software, which extend its functional and security capabilities. This chapter deals primarily with English auction software mechanisms (see Section 2.1). Chapter 6 introduces software bidding agents and Chapter 8 describes CDA related software.

2.1 RAS: A System for Supporting Research in Online Auctions

This paper describes our experiences with designing online auction software and presents a client-server software model for conducting online English auctions. The model is based on an existing auction server developed at James Cook University called the "*Research Auction Server*" (*RAS*). *RAS* is used to perform both simulated and real auctions to gather data on the performance of key areas of auction research. *RAS* has been used to conduct research in auction security, fraudulent bidding behaviour, bidding agent design and market clearing algorithms. This paper discusses the mechanics of English auctions, and describes the major software components of an online auction and the processes involved. It also addresses website navigation, object-oriented design, database construction, transaction issues, and timing concerns.

"*RAS: A System for Supporting Research in Online Auctions*" [6] is published in ACM Crossroads. Crossroads is a peer-reviewed student journal by the Association for Computing Machinery (ACM). Crossroads is published quarterly and is provided to ACM members as part of their subscription. Each issue has a theme (e.g., programming languages, artificial intelligence, etc.). The theme for this issue is *Software Engineering*. The review process involves critical analysis from at least two Associate Editors as well as other ACM staff. An accepted paper undergoes a rigorous shepherding phase where it is guided to completion by a further Associate Editor. Only high quality papers with a significant contribution to their field are accepted for publication. This paper is available online via the Crossroads website ¹, ACM Portal ², citeseer ³, and is referenced by DBLP ⁴.

¹<http://www.acm.org/crossroads/>

²<http://www.portal.acm.org>

³<http://citeseer.ist.psu.edu/>

⁴<http://www.informatik.uni-trier.de/~ley/db/>

The main requirements for RAS and preliminary prototype were developed in 2003 as part of a research project where I supervised third year IT student Colin Ellems. RAS has been made possible by hardware and software grants from the School of Mathematical and Physical Sciences in 2005 and the combined School of Mathematics, Physics and Information Technology in 2006. In addition, Zac Burrell provided some initial technical support.

RAS: A System for Supporting Research in Online Auctions

by Jarrod Trevathan and Wayne Read



Introduction

Online auctioning is unparalleled as the fastest growing exchange medium to emerge from electronic commerce technology. Buyers and sellers located around the world now auction various items from the latest DVD to rare collectibles. eBay [2] and uBid [10] are among the most successful and popular of the commercial online auctioneers. They use an auctioning process based on a type of auction referred to as an *English* auction. In an English auction, bidders outbid each other for an item. The winner is the bidder with the highest bid.

Despite the popularity of auctioning online, there are inherent security risks. For example, bidders might not pay or sellers might not deliver the item. To compound the problem, commercial auctioneers do not hold themselves liable for any transactions that go awry. Cryptographic and other security methods can help protect against many of these risks. However, there is no way to test the feasibility of such auction mechanisms. Therefore, an evaluation environment for research purposes is required.

In 1998, Wellman et al. [11] developed an online auction server called the AuctionBot. Designed at the University of Michigan, the AuctionBot's main purpose was to serve as an infrastructure for conducting auction research. The AuctionBot strived to be *generic* by its capacity to run various types of auctions (i.e., English, Vickrey, and so on). Specific types of auctions could be performed by altering the AuctionBot's parameters. The AuctionBot also contained an agent interface for software bidding agents.

Software bidding agents bid on behalf of a human bidder according to a predetermined strategy. After the AuctionBot was decommissioned in 1999, its platform became the basis for the 2001 Trading Agent competition [14]. The competition pitted bidding agents against each other. The agents participate in an elaborate economic game and are assessed on their ability to acquire certain resources. Other literature on the AuctionBot can be found in [12,13].

In 1999, Kumar and Feldman [4] presented a software model for auctioning online. The design was based on an auction system they had implemented. Similar to the AuctionBot, they also sought a generic auction architecture. Sound software engineering principles are applied by employing an

object-oriented approach to auction design. Kumar and Feldman describe auction web site management, navigational issues and the process flow of an auction. In addition, they discuss the effect of timing and security issues with online auctions.

Despite these initial attempts at auction software design, publicly available research activity has since receded. This is largely due to the dominance of commercial online auctioneers. To compound the problem, commercial auctioneers tend not to publish their research.

We required our own infrastructure for performing research. There are many vendors selling auction software online. However, such software is expensive and cannot be customized for our research needs. Due to this fact and the limited availability of academic literature on auction software design, we have created our own online auction server.

This article describes our experiences with designing online auction software and presents a client-server software model for conducting online English auctions. The model is based on an existing auction server developed at James Cook University called the "*Research Auction Server*" (RAS). We use RAS to perform both simulated and real auctions and to gather data on the performance of key areas of auction research. RAS has been used to conduct research in auction security, fraudulent bidding behavior, bidding agent design, and market clearing algorithms. RAS is open source and available online at <http://auction.maths.jcu.edu.au>.

In this article we discuss the mechanics of English auctions, then we describe the major software components of an online auction, the processes involved, and web site navigation. Next, we present an object model of the system along with a database schema and transaction issues. We address timing concerns and specify a software bidding agent interface. Finally, we discuss research conducted on RAS and its implications for transaction settlement.

Online English Auctions

There are many types of auctions such as English, Vickrey, Dutch, double, and so on. (See Trevathan et al. [8].) The *English* auction is the most well known type of auction.

Formally, an English auction is referred to as an *ascending price* auction. To win, bidders must outbid each other, forcing the price up. The winner is the bidder with the highest price at the end of the auction.

The form of a bid in an English auction is:

< bidder, price, time >

At any time in an online auction, a bidder can request a *price quote* from the auctioneer. The price quote contains the bid information of the current highest bid.

The *bid history* of an auction is a list of all bids that have been submitted. Table 1 gives an example bid history for an auction. Here the winning bid is \$55.

Bid	Bidder	Price	Time
7	Wayne	\$ 55	33:05
6	Jarrold	\$ 50	30:15
5	Sharith	\$ 40	24:22
4	Wayne	\$ 35	17:01
3	Mr Gab	\$ 25	12:35
2	Shane	\$ 20	05:44
1	Jarrold	\$ 15	02:09

Table 1: Bidding History.

English auctions can have three additional parameters:

- *Starting price*—the minimum price at which the bidding must commence.
- *Reserve price*—the minimum price the seller of the item will accept. If the final price is below the reserve, the auction result is void.
- Minimum *bid increment*—the minimal amount required to outbid the current highest offer.

Online auctions allow many variations on English auctions. For example, bidders can alter or cancel bids, auction multiple goods, and issue sell bids. We will not deal with all of the possible extensions to English auctions, but rather restrict our attention to the essential features.

Previous auction software architectures have strived to implement as many auction types as possible. However, the model in this article only focuses on English auctions. There are several reasons for this. First of all, English auctions are

the most common type of auction performed online. Secondly, there is limited literature on software design for online English auctions. Furthermore, some auctions (for example, continuous double auctions) are more complicated than English auctions. Therefore our auction software model allows a sufficient amount of detail.

Components of an Online Auction

Figure 1 presents a high level software model for performing online English auctions. There are two main parties: a bidder and an auctioneer. The parties are joined by a communication link.




Figure 1: Online English auction software model.

There are two types of interfaces for a bidder. The first is the web interface. This is for a human bidder. The bidder interacts with the auctioneer via an HTML browser. The second interface is for a software bidding agent. A bidding agent interacts with the auctioneer using an application programming interface.

The auctioneer runs a web server (e.g., Apache) and a scripting language, in this case PHP. The entire auction is database driven. All state information (e.g., bids and timing) about the auction is contained in the database. When a client submits a bid or requests a price quote, a database transaction occurs. The database generates dynamic web pages in response to bidder activity using the scripting language.

Processes Involved

There are several main activities in an online English auction:

- *Initialization*: The auctioneer sets up the auction and advertises it, i.e., type of good, starting time, etc.
- *Registration*: In order to participate in the auction, bidders must first register with the auctioneer.
- *Price quote*: A bidder obtains a price quote from the auctioneer.

- *Bidding*: A bidder submits a bid to the auctioneer.
- *Winner Determination*: The auctioneer determines the winner according to the auction rules.

Transaction settlement/payment is the process of collecting payment from and delivering the goods to the winning bidder.

Bidder/Winner Notification refers to the processes involved with informing a bidder with information other than the price quote. This is information specific to an individual such as confirmation of bid receipt or notifying the winner that they have won.

Web Interface Navigation

Figure 2 illustrates how human bidders navigate the auction Web site. Each bubble represents a major section of the site. Lines represent links between sections and arrows indicate the navigational direction a user must follow to reach a particular page.

After registering, a user can log in using a password. Upon login a user enters the secure area (shown by the red box in Figure 2). This is a collection of pages which contains operations that only a registered user can perform.

The home page is the main area for the bidder. Bidders are able to search for a listed auction. Once a bidder has selected an auction, they are able to obtain a description of the item

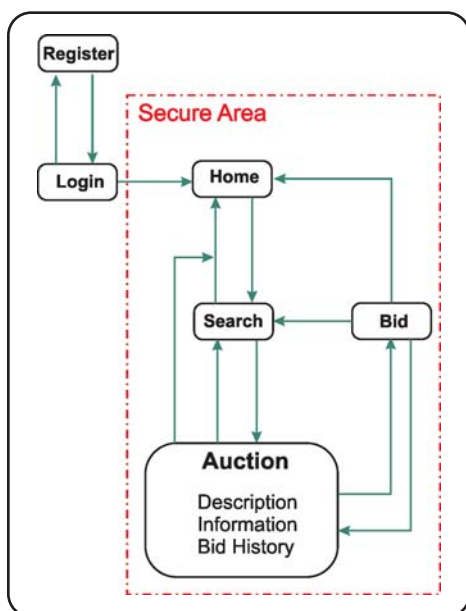


Figure 2: Web interface navigation map.

and information regarding the auction (i.e., price quote, minimum increment, and time remaining). The bidder can then use this information to submit a new bid. After submitting a bid, a bidder can return to the home page, search for a new auction, or return to the current auction.

In the secure area, specific information regarding the user (such as user id and password) must be carried from one page to another. To achieve this, the secure area is implemented using session variables. Session variables are similar to cookies, as they are used to store information for a particular period of time. The values in session variables exist until the session terminates. This allows information to be passed between web pages or to another web site. A session can be created using a session identifier which is stored on the server. When the client makes a request, the data stored in the session variable can be accessed any number of times until the session ends.

An alternative approach for passing information between pages is to use a query string. However, using a query string in this manner is difficult and cumbersome. Session variables on the other hand allow information to be more easily retrieved and maintained. Figure 3 illustrates the process involved for using session variables.




Figure 3: Session variable authentication procedure.

When a bidder logs off, the session variable is destroyed. The variable is also destroyed after a predetermined amount of bidder inactivity (for example, when a user forgets to logoff).

Upon submitting a bid, the bidder receives an email confirming that his/her order has been received. The bidder will also be informed of such information via the web page.

Search is an important mechanism in online auctions. Bidders have the ability to choose from millions of auctions to

participate in. The search mechanism must allow a bidder to quickly and accurately identify a desired auction. Items for auction are typically categorized hierarchically. For example, the category “automobiles” would contain auctions for cars, motorbikes, trucks, etc. The “car” category would be further decomposed into auctions for sedans, hatchbacks, station wagons, etc. Other items for auction might be listed according to the seller. A full description of searching methods is beyond the scope of this article, however, we leave the architecture of RAS open to incorporate any search mechanism.

Object Model

The web interface for RAS allows human bidders to interact with the auctioneer. Figure 4 depicts an object model for RAS’s web interface. There are two objects: *Bidder* and *Auctioneer*. The top portion of each object indicates its internal state and the bottom portion lists its methods (see Coad and Yourdon [1]).




Figure 4: Web interface object model.

A *Bidder*’s internal state is the *userId* and *password*. The bidder interacts with the auctioneer via the auctioneer’s methods. The bidder only has two methods: *register* and *determine bid*. During registration, a bidder obtains a *userId* and *password*, which are kept secret. When participating in an auction, a human bidder determines his/her bid according to individual preferences.

The *Auctioneer* manages the database, which is the central component of the auction. This contains all information regarding bidders, their bids, and the auctions conducted. There are two types of methods the auctioneer can perform: private and public. Private methods are operations that only the auctioneer can perform, such as *list a new auction* (*list auction*), *register* new bidders, and *terminate* an auction. Public methods are services which are requested by a bidder. This includes the ability to *login*, *search* for an auction, request the minimum *increment*, obtain a *price quote*, request the *time remaining*, submit a bid (*submit bid*), and *logoff*.

Database

Figure 5 depicts an entity-relationship diagram for the database in an online English auction. Rectangles represent enti-

ties, diamonds represent relationships, and ellipses represent attributes (see Silberschatz et al. [5]). There are three entities: *User*, *Auction*, and *Bid*. The primary key for each entity is underlined.

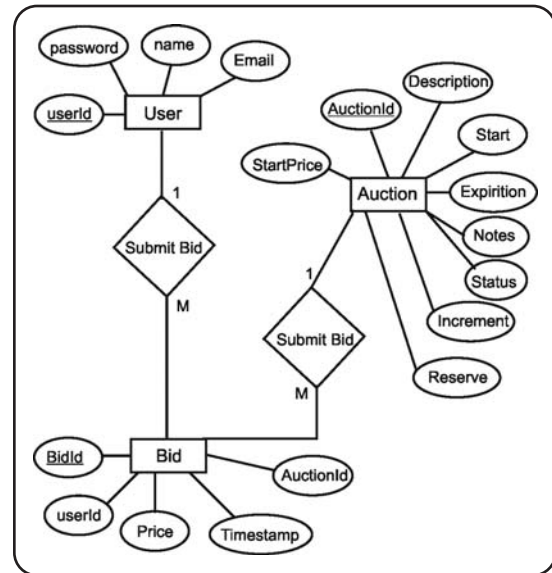


Figure 5: Entity relationship diagram.

When a bidder registers, their details are entered into the user entity. Each user is identified by a unique *userId*. Information stored about a user includes their *password*, *name*, and *email* address.

Each auction is identified by a unique *auctionId*. Information stored about an auction includes an item *description*, *start* time, *expiration* time, *startPrice*, *reserve* price, minimal *increment*, and *notes* about the item. The *status* attribute indicates whether the auction is currently active or has terminated.

Every time a bidder submits a bid, all information about the bid is entered into the bid entity. Each bid is identified by a unique *bidId*. To associate a bid with a bidder, the bid entity stores the *userId* as a foreign key. To associate a bid with an auction, the bid entity also stores the *auctionId* as a foreign key. Other bid information includes the *price* of a bid and a *timestamp* that indicates when the bid was submitted.

In this model a user cannot list items in an auction directly. Instead this must be done exclusively by the auctioneer. Allowing a bidder to list items for sale requires an extension to the entity-relationship diagram. The auction essentially becomes a double auction (i.e., many buyers and many sellers), which complicates matters.

There are several database transactions that occur during an auction:

- Initialization—A new auction’s details are entered into the database.
- Registration—A new user’s details are entered into the database.
- Login—An existing user’s details are retrieved from the database.
- Price Quote—An existing auction’s details and bid history are retrieved from the database.
- Bid Submission—A new bid’s details are entered into the database.
- Auction Termination/Winner Determination—An existing auction’s details and bid history are retrieved from the database. The auction’s status is set to ‘inactive’ and entered in the database. The winner is determined according to the auction’s rules.

Timing

When a bid is received for a particular auction, it is compared to the current highest bid. If it does not meet the minimal amount required to outbid the existing highest bid, then it is discarded. A bid that does meet the required amount is timestamped by the auctioneer and entered into the database.

Online English auctions can terminate according to the following rules (see [4, 6]):

- Expiration Time—The auction closes at a predetermined expiration time.
- Timeout—The auction closes when no bids higher than the current highest bid are made within a predetermined timeout interval.
- Combination of Expiration and Timeout—The auction closes when there is a timeout after the expiration time.

In our model, the auction is database driven. Rather than having a running program continually check the time in order to terminate the auction, our auction terminates in response to bidder activity. This means that every time a bidder submits a bid or requests a price quote, a script is

run to check whether the auction has terminated according to the database clock. This involves comparing the database time to the expiration attribute of the auction entity in the auction entity-relationship diagram (Figure 5). A pseudo-code script for terminating an auction with an expiration time is shown below (in blue):

```

if database time >= expiration {
  set auction status to inactive
  select highest bid from auction
  if bid > reserve
    declare this bid the winner
  else
    auction is void
}
else {
  expiration = database time + timeout
}

```

This setup does not require synchronization between the bidders and the auctioneer.

Terminating an auction using a timeout involves adding additional time to the auction’s expiration attribute in the database every time a higher bid is received. Let *timeout* be a constant time increment. Adding a timeout extends the terminate auction script as shown by the red section of the code.

Some variants of English auctions allow bid cancellation. This is common in English auctions which terminate using an expiration time. The justification for this is that such auctions can often take days or weeks. In this situation, a bidder may be reluctant to make such an open-ended bid.

The simplest approach to bid cancellation is to delete the cancelled bid from the database. However, for security reasons a cancelled bid should be recorded. This is achieved by including *valid* and *cancellationTime* attributes in the bid entity of the entity-relationship diagram shown in Figure 5. Figure 6 illustrates new additions to the bid entity where the red indicates the new attributes.




Figure 6: Bid entity with bid cancellation.

A bid is initially in a valid state and the cancellation time is *null*. This indicates that it is to be considered for an auction. If a bid is cancelled, it enters an invalid state, indicating that it should not be considered for an auction. This allows the bid details and time of cancellation to be recorded but not included in an auction.

Bidding Agents

A bidding agent is a program which bids on behalf of a human bidder. In terms of an English auction, a bidding agent is permitted to outbid any bid until the bidding price exceeds a maximum amount specified by the human bidder. In auctions that can last days or weeks, bidding agents remove the need for a bidder to constantly observe an auction. Instead, a bidding agent monitors the auction proceedings for any price activity and responds in accordance with its programmed strategy.

The object model for a simple bidding agent is shown in Figure 7. The auctioneer must provide an agent with six basic services: *login*, *request price quote*, *minimum increment*, *time remaining*, *submit bid* and *logout*. We have not listed the auctioneer's private methods in the model.




Figure 7: Simple bidding agent object model.

Similar to a human bidder, an *Agent* object has a *userId* and *password*. An agent also knows the *auctionId* of the auction it wants to participate in. The *valuation* is the maximum limit that an agent can bid. An agent can perform two operations: *initialize* and *determine bid*. The *initialize* method provides the agent with the human's *userId* and *password*, and instructs the agent which auction to participate in and the valuation of the auctioned goods.

The following pseudocode script of the *determine bid* method illustrates the basic operation of a bidding agent. Green code indicates the methods of the Auctioneer object that are used by the agent.

```

bid = 0
login ( userId, password, auctionId )
  
```

```

repeat {
  if ( bid <= price quote ) {
    bid = price quote + increment
    if ( bid <= valuation AND time remaining )
      submit bid ( bid )
    else
      logout
  }
}
  
```

Initially the agent's bid is zero. The agent logs in by supplying its *userId* and *password*. The agent also indicates the auction it wants to participate in by specifying the *auctionId*. Once logged in, an agent repeatedly requests a price quote and alters its strategy accordingly.

In our simple agent example, the agent's current bid is compared to the current highest bid in the auction. If the agent's bid is less than the price quote, then the agent increases its bid by a given increment. If the new bid does not exceed the agent's preset valuation, then the agent submits the bid to the auctioneer. Otherwise the agent logs off as it is not allowed to exceed the valuation. Before submitting a bid, the agent also checks whether the auction has terminated. If so, the agent logs off.

Bidding agents are not limited to the strategy employed by our simple agent example. Bidding agents can employ sophisticated strategies. Some of these strategies are based on statistical analysis of the auction data and neural network reasoning about the best strategy to take.

Research Conducted Using RAS

This section briefly describes research we have been conducting on RAS.

Security and Anonymity

Security and anonymity are crucial in online auction software design. Unlike a traditional auction, the participants in an online auction are not physically present. This presents many opportunities for cheating. For example, a bidder might repudiate having made a bid, forge a bid on behalf of another bidder, or refuse to pay. Furthermore, the auctioneer might be corrupt and award the auction to someone other than the legitimate bidder. Additionally, outsiders may attempt to influence or disrupt the auction proceedings. Online auctioning also raises many concerns regarding the privacy of a bidder's personal information, e.g., identity and bidding history.

Security in auctions has been discussed in the literature (for example, see [4, 6, 7]). Moreover, various cryptographic auctioning schemes have been proposed. The main goals of such schemes are to prevent the forging of bids and bid repudiation, to protect against auctioneer corruption, and to preserve the anonymity of bidders. Cryptographic schemes attempt to achieve this by signing and encrypting bids, and by employing anonymity preserving mechanisms and publicly verifiable protocols.

A secure scheme must be able to provide security with a reasonable quality of service to auctions involving large numbers of bidders. Cryptographic operations such as bid signing/verification and encryption are computationally expensive, requiring time to perform. This effectively slows down the auctioning process and may result in bids not being included in an auction.

RAS is used as a platform to develop and test security mechanisms for online auctions. RAS has been used to implement existing cryptographic auction schemes in order to determine their feasibility by measuring signing/verification and encryption times. We have also proposed schemes of our own (see [8, 9]), which we have implemented using RAS.

RAS is also used to develop mechanisms to detect and prevent fraudulent bidding practices such as shilling. Shilling is where the seller introduces fake bids to drive up the final price. This is a problem as it forces a legitimate bidder to pay more for an item. English auctions are particularly susceptible to shill bidding behavior. We have identified the core strategies of a shill bidder and are developing methods to detect shill bidding. Both real and simulated auctions are conducted, some of which involve fraudulent bidding behavior. This allows us to test the effectiveness of our fraud detection techniques and to observe the reactions of innocent bidders.

Software bidding agents are used to simulate auctions involving large numbers (i.e., thousands) of bidders.

Transaction Settlement/Payment Mechanism

The payment mechanism is used to settle the auction (i.e., collect payment from the winner and ensure delivery of the item bid upon). The method employed by RAS depends on whether the auction is simulated or real.

During simulated auctions, fake items are auctioned. Bidders are issued a fake amount of money to use. At the end of the auction, the winning bidder does not receive the item nor

are they required to pay. Such auctions are designed to emulate real auctions without concern over payment settlement.

Real auctions, on the other hand, involve auctioning real, tangible items. Bidders are also required to use their own money. When the auction finishes, the winner receives the item bid upon and must also pay the auctioneer the winning amount. Such auctions are used to gauge actual auctioning behavior, which can differ from simulated behavior as there are real risks and rewards.

Enforcing payment in real auctions is a complicated task. This often involves setting up accounts, providing insurance, gathering user feedback on sellers, etc. We choose not to tackle this topic in this article. Instead we leave the architecture open to implement any payment mechanism for real auctions.

Simulated auctions provide the system with a great deal of control. Each user can be issued an initial balance. When a bidder wins an auction, the amount of the winning bid is deducted from the bidder's balance. A user is unable to submit a bid for an amount greater than their balance.




Figure 8: User entity with payment mechanism.

The user entity in the entity-relationship diagram in Figure 5 can easily be modified to include this behavior. Figure 8 illustrates the amended user entity. The red section depicts a new attribute called *balance*. The balance attribute contains the current balance of the user. When a bidder wins an auction, his/her balance attribute is updated.

Conclusions

Our auction model is specific to online English auctions and it is simple and concise compared to previous literature on auction design. Our system is modeled using an object-oriented approach and is used to test the efficiency and practicality of security and anonymity mechanisms in online auctions. RAS has been used to perform hundreds of simulated and real auctions.

RAS is being extended to other types of auctions. For example, we have implemented a form of auction referred to as a

continuous double auction (CDA). This allows many buyers and sellers to continuously trade goods. The best known example of a CDA is a share market. A CDA is inherently more sophisticated than an English auction. This involves modifications to the database schema and the winner determination procedure. We are in the process of documenting the software design considerations of RAS's CDA mechanism.

Acknowledgments

The authors would like to thank Dr. Sharith Trevathan for her contribution to this article.

References

- 1 Coad, P. and Yourdon, E. *Object-Oriented Analysis*, 2nd Ed. Yourdon Press Computing Series, Prentice-Hall PTR (1991).
- 2 eBay at <http://www.ebay.com>.
- 3 Franklin, M. and Reiter, M. The design and implementation of a secure auction service. *IEEE Trans. Softw. Engin.* 22 (May 1996) 302-312.
- 4 Kumar, M. and Feldman, S. I. Internet auctions. In *Proceedings of the 3rd USENIX Workshop on Electronic Commerce* (Aug. 1998). 49-60.
- 5 Silberschatz, A., Korth, H., and Sudarshan, S. *Database System Concepts*. McGraw Hill Computer Science Series (1996).
- 6 Stubblebine, S., and Syverson, P. Fair online auctions without special trusted parties. In *Proceedings of Financial Cryptography*. M. Franklin, Ed. Lecture Notes in Computer Science, vol. 1648, Springer-Verlag (1999). 230-240.
- 7 Trevathan, J. Security, anonymity and trust in electronic auctions. *ACM Crossroads* 11, 3 (Spring 2005) 3-9.
- 8 Trevathan, J., Ghodosi, H., and Read, W. Design issues for electronic auctions. In *2nd International Conference on E-Business and Telecommunication Networks (ICETE'05)* (Oct. 2005). 340-347.
- 9 Trevathan, J., Ghodosi, H., and Read, W. An anonymous and secure continuous double auction scheme. In *39th Hawaii International Conference on System Sciences (HICSS'06)* (Jan. 2006). 125.
- 10 uBid at <http://www.ubid.com>.
- 11 Wellman, M. P., Wurman, P. R., and Walsh, W. E. The Michigan Internet AuctionBot: A configurable auction server for human and software agents. In *2nd International Conference on Autonomous Agents (AGENTS)* (May 1998). 301-308.
- 12 Wellman, M. P., and Wurman, P. R. Real-time issues for Internet auctions. In *1st IEEE Workshop on Dependable and Real-Time E-Commerce Systems (DARE'98)* (1998).
- 13 Wellman, M. P. and Wurman, P. R. Parameterization of the auction design space. *Games Econom. Behav.* 35 (2000) 304-338.
- 14 Wellman, M. P. and Wurman, P. R. The 2001 trading agent competition. *Electronic Markets* 13, 1 (2003) 4-12.

Biography

Jarrod Trevathan (jarrod@cs.jcu.edu.au) is a PhD student at the School of Mathematical and Physical Sciences, James Cook University. His research interests include security, programming, multimedia, and electronic commerce. Jarrod has also taught Computer Science for six years.

Wayne Read (wayne.read@jcu.edu.au) is Head of School at the School of Mathematical and Physical Sciences, James Cook University. His research interests include modelling transport processes in porous media, series solutions to partial differential equations, and restricted occupancy models.



Be Remarkable.

BAE Systems is a high-technology aerospace and defense engineering company which provides intelligent electronic systems, information technology and technical services for homeland security, defense and commercial aerospace customers.

Advanced Information Technologies (AIT), is the Applied Research arm of BAE Systems located in Burlington, MA

We want you to bring your advanced degree(s), ambitious ideas and cutting edge research to us. You will have the opportunity to work on some of the most important problems facing our country today in both the Defense and Intelligence communities.

Applicants selected may be subject to a security investigation and must meet eligibility requirements for access to classified information. For a career beyond all expectations, please visit our website. EOE, M/F/D/V.

www.ait.na.baesystems.com

BAE SYSTEMS

Chapter 3

Auction Privacy and Security Issues

This chapter presents two papers that raise awareness of online auction privacy and security issues. These papers set the theme for the rest of this thesis. The first paper introduces the privacy and security problems inherent in the auctioning process (see Section 3.1). The second paper investigates existing cryptographic solutions and identifies several flaws in current security design proposals (see Section 3.2). These papers describe the principles underlying the proposed auction security and anonymity preservation techniques.

3.1 Security, Anonymity and Trust in Electronic Auctions

This paper discusses security, anonymity and trust issues in electronic auctions. It outlines exactly how bidders can cheat in existing commercial online auctions. In addition, it shows how the Auctioneer can influence the auction in a manner that is inconsistent with the rules. Existing auction security schemes use various Auctioneer configurations to ensure that the Auctioneer can be trusted. The paper contrasts these configurations describing their advantages and disadvantages. Concerns regarding abuse of a bidder's identity and personal information are addressed. Furthermore, this paper describes how an outsider can influence the auction proceedings by forging bids and through denial of service attacks. Various types of fraudulent bidding behaviour are also discussed.

“Security, Anonymity and Trust in Electronic Auctions” [2] is published in ACM Crossroads. Crossroads is a peer-reviewed student journal by the Association for Computing Machinery (ACM). Crossroads is published quarterly and is provided to ACM members as part of their subscription. Each issue has a theme (e.g., programming languages, artificial intelligence, etc.). The theme for this issue is *Cryptography*. The review process involves critical analysis from at least two Associate Editors as well as other ACM staff. An accepted paper undergoes a rigorous shepherding phase where it is guided to completion by a further Associate Editor. Only high quality papers with a significant contribution to their field are accepted for publication. This paper is available online via the Crossroads website ¹, ACM Portal ², citeseer ³, and is referenced by DBLP ⁴.

¹<http://www.acm.org/crossroads/>

²<http://www.portal.acm.org>

³<http://citeseer.ist.psu.edu/>

⁴<http://www.informatik.uni-trier.de/~ley/db/>

As a result of having two publications in *Crossroads*, I have now held the position of Associate Editor since January 2006. In 2004 I traveled to Key West Florida USA to attend the Financial Cryptography conference (FC 2004). This paper is the culmination of my initial research into auction security. Travel to the conference was made possible by the Doctoral Research Scheme (DRS) and the School of Information Technology. This paper resulted in a media article by the *Townsville Bulletin*.

Security, Anonymity and Trust in Electronic Auctions



by Jarrod Trevathan

Introduction

Auctioning items over the Internet is a popular and lucrative industry. There are now many companies that conduct auctions online such as *eBay* [5] and *onSale* [10]. Online auctions have geographical advantages over traditional auctions as buyers and sellers are not required to be physically present at a central location (such as a hall or open air venue). This allows online auctions to be much larger and more elaborate than traditional auctions. However, it also provides opportunities for the auction participants to cheat.

A bidder can cheat by repudiating bids, failing to pay, or colluding with other bidders to affect the settlement price. Likewise, the seller of the item might fail to deliver the goods, or could be in collusion with some of the bidders. Someone could also forge a bid in an attempt to frame a bidder, or introduce fake bids in order to influence the auction proceedings.

Furthermore, bidders are required to trust the auctioneer with their identity and bid information. A corrupt auctioneer could award the auction to someone other than the legitimate winner. A bidder's personal information could also be sold to marketing agencies, or used for malicious purposes.

Commercial auction sites fail in many of the aforementioned circumstances. These sites only offer basic solutions that are designed to "clean up" after wrongdoing has taken place. However, cryptography can be used to solve some of these problems up-front. An "electronic auction" is a cryptographic scheme designed to securely conduct auctions while protecting the identities of the bidders.

In this article we describe two popular types of electronic auctions. We discuss the security issues associated with conducting these auctions and contrast the differing anonymity requirements. We also identify four main strategies for reducing the trust that bidders must place in the auctioneer. Furthermore, we present a basic example of an electronic auction scheme. This is used to illustrate the complexity involved in designing a secure and anonymous auction scheme. Finally, we discuss some of our research with regard to using group signature schemes to construct electronic auctions.

Types of Auctions

The first known auctions were conducted in Babylon circa 500 BC with the dubious application of auctioning women for marriage. The idea was extended to various goods and services in Ancient Rome [4]. There are now a variety of auction types that have evolved to suit differing needs and applications.

The most well known auction is the *English auction*. This type of auction is commonly used in Real Estate where a bidder attempts to outbid all other bidders. The winner is the bidder with the highest bid and they must pay an amount equal to this bid. An English auction is referred to as an *open bid auction* because everyone knows the amount that a bidder has bid.

A *Vickrey auction* is a lesser-known auction where bid amounts remain *sealed*, meaning no one knows the value of anyone else's bid. Vickrey auctions were invented in 1961 by Nobel Prize winner William Vickrey [14]. A Vickrey auction consists of two stages: bidding and opening. During the bidding stage, bidders seal their bid (e.g., place it in an envelope) and submit it to the auctioneer. During the opening stage, the auctioneer opens all of the bids and determines the winner. The winner is the bidder with the highest bid. They are required to pay an amount equal to the second highest bid (i.e., the highest losing bid).

Vickrey auctions are popular in cryptographic research as they are less complicated to model compared to English auctions. This is because English auctions require a real-time communication channel between the auctioneer and bidders for price updates and winner determination. Vickrey auctions, on the other hand, only require a bidder to send a single message to the auctioneer.

Security

Online auctions are inherently insecure due to geographical scale and lack of accountability. Since participants are not physically present at the auction proceedings there are opportunities to cheat. For example, bidders can repudiate or forge bids and sellers might not deliver goods. Furthermore, the auctioneer could be corrupt. Auction security has been studied extensively in literature [1, 2, 6, 7, 8, 9, and 13]. The primary goal of an electronic auction scheme is to ensure that the auction outcome is computed fairly. To understand

the issues involved, we first review some problems encountered by traditional auctions and then examine various solutions employed by existing commercial auction sites.

Some problems with traditional English and Vickrey auctions include:

- **Shielding:** A bidder places an artificially high bid which is subsequently withdrawn just prior to the bid close time. This can have the effect of deterring other bidders with lower valuations from bidding. When the bid is withdrawn, the bidder may have another lower bid ready to win at the deflated price.
- **Shilling:** A bidder, or group of bidders (called shills), collude to artificially inflate the clearing price for the seller. If one of the shills accidentally wins, the item is resold in another auction.
- **Sniping:** A bidder refrains from making a bid until just prior to the bid close time. When the bid is made it is usually done so in a manner that does not allow any other bidder to respond in time.
- **Siphoning:** A non-participant observing an auction makes a lower offer directly to a bidder. The non-participant avoids the costs and risks associated with conducting an auction.
- **Misrepresented or non-existent items:** A seller might make false claims about the item for sale, or attempt to sell an item they do not have.

In an attempt to reduce these problems, existing commercial online auction sites offer remedies such as legislation and incentives. Firstly, laws can be made with regard to breaking the rules. For example, defaulting on payment could result in a fine and/or jail sentence. Another solution is to provide an incentive scheme to reward people that follow the rules. For example, eBay [5] offers a feedback system that allows buyers and sellers to create profiles about each other based on previous dealings. Anyone can view these profiles before engaging in business with the individual. A buyer or seller with a shady reputation can also be blacklisted from future auction proceedings.

Other solutions involve credit card registration, escrow services and insurance. Credit card registration can be used as

both proof of identity and security for payment. In a similar manner, escrow services require all bidders to keep payment in escrow (security), which is either refunded or deposited depending on the auction outcome. Alternately, insurance can be offered to participants in situations where they suffer loss as a result of unfair behavior. However, none of these solutions is perfect.

Electronic auctions seek to use cryptographic protocols to enforce security and protect anonymity. This is done by constructing auction schemes from cryptographic building blocks such as encryption, digital signatures, secret sharing, and digital cash. First the basic components of an electronic auction scheme must be described.

There are several main activities (or stages) fundamental to an electronic auction protocol:

- **Initialization:** The auctioneer sets up the auction and advertises it (i.e., type of good, starting time, etc.).
- **Registration:** In order to participate in the auction, bidders must first register with the auctioneer (or a registration manager). This ensures that only valid bids are made and that bidders can be identified for payment purposes. It is desirable for registration to be a one-off procedure. Once a bidder has registered they should be able to participate in any number of auctions rather than re-registering for each new auction.
- **Bidding:** A registered bidder computes his/her bid and submits it to the auctioneer. The auctioneer checks the received bid to ensure that it conforms to the auction rules.
- **Winner Determination:** The auctioneer determines the winner according to the auction rules. It is desirable for this process to be publicly verifiable.

The following outlines the main security goals for electronic auction schemes:

- **Un-forgable bids:** Bids must be un-forgable, otherwise a bidder can be impersonated.
- **Non-repudiation:** Once a bidder has submitted a bid they must not be able to repudiate having made it. For example, if a bidder wins and does not want to pay, they might deny that they submitted the bid.

- **Publicly verifiable:** There must be some publicly available information by which all parties can be verified as having correctly followed the auction protocol. This should include evidence of registration, bidding, and proof of winner/loser.
- **Robustness:** The auction process must not be affected by invalid bids or by participants not correctly following the auction protocol.
- **Efficiency:** Efficiency issues also play a large role in the practicality of electronic auction schemes. Factors that influence the efficiency of a scheme include the computational and communication overhead required for registration, signing a bid, verifying a bid and winner determination.

Many of the schemes proposed in literature fail to achieve the basic security goals or are too inefficient to be practical. Furthermore, some schemes reveal too much information about a bidder, including his/her identity and bidding history.

Anonymity

A major problem with existing commercial auction sites is that there is no means to protect a bidder's identity. For example, the auctioneer knows everything about a bidder including his/her identity and the values of the bids he/she submits. This is undesirable as such information can be used to target a bidder with unsolicited junk mail or for more malicious purposes such as bid shielding.

There are varying levels of anonymity depending on whether the auction is English or Vickrey. However, the main objective for anonymity in an electronic auction scheme is:

To conceal the bidder-bid relationship so that no bidder can be associated with the bid they submit.

Anonymity is often achieved by issuing bidders with a pseudonym (fake identity) that they can use to submit bids. Consider the following scenario for example: A bidder must register with two registration managers, denoted as RM1 and RM2 respectively. RM1 knows the bidder's real identity and issues the bidder with a certificate. The bidder presents this certificate to RM2. RM2 trusts the certificate and provides the bidder with a pseudonym that can be used when bidding. RM1 knows the bidder's identity, but not the corresponding pseudonym. RM2 knows the bidder's pseudonym, but not his/her identity. Here the bidder remains anonymous as long as RM1 and RM2 do not collude.

However, anonymity is at odds with non-repudiation. If an auction were completely anonymous, bidders would be able to repudiate bids as there is no way to tell who submitted what bid. Therefore, it is desirable to have a mechanism to trace bidders in the event of a dispute. This can be thought of as an "identity escrow" scheme where an authority has the power to reveal the identity of who submitted the bid in question. In the example above, RM1 and RM2 could be called upon, under special circumstances, to reveal a bidder's identity by combining their information.

Further anonymity issues include:

- **Confidentiality:** In a Vickrey auction once a bid has been sealed, it must remain sealed until the winner determination stage of the auction.
- **Privacy of losing bids:** In a Vickrey auction, the values of the losing bids must be kept secret (however, most schemes leak bid statistics). Further questions also arise over how much information is disclosed regarding the winner and winning bid. For example, should only the seller and winner know the amount of the winning bid? What information should the auctioneer learn about winning/losing bids?
- **Un-linkable bidding:** The auctioneer should not be able to learn information about individual bidders based on previous auctions conducted. This information could be used in future auctions in a manner that disadvantages the bidder. Consider the following scenario in a Vickrey auction: a bidder (denoted by B1) submits a bid for \$100. The second highest bid is \$50, so B1 wins and has to pay only \$50. In a future auction, B1 again bids \$100 and again the second highest bid is \$50. Since the auctioneer knows B1's valuation, the auctioneer enters a bid for \$99. B1 wins and must pay \$99. In doing so, the auctioneer has increased the auction's revenue by \$49.

In an English auction [9], linkable bidding within an auction is acceptable, as it is common for several bids to come from one particular bidder. For example, two bidders might engage in a rally attempting to outbid each other. Observers and participants of an auction often gain entertainment value from watching the parties involved in such rallies. However, linking bids between subsequent auctions is not desirable. An auctioneer who knows the valuation of a bidder might set a higher reserve (minimum winning) price in future auctions involving that bidder.

The need for anonymity in a scheme must be weighed against the goals and purpose of the scheme. Too much anonymity allows bidders to repudiate bids, whereas not enough anonymity allows bidders to be profiled and cheated by the auctioneer. To help with issues of security and anonymity, it is therefore desirable to reduce the level of trust a bidder must place in the auctioneer.

Trust

A significant problem in electronic auction protocols is how to protect bidders from a corrupt auctioneer. A malicious auctioneer might influence the auction proceedings in a manner inconsistent with the auction rules. For example, the auctioneer might choose to block bids, insert fake bids, steal payments, profile bidders, open sealed bids prior to the opening stage, or award the item to someone other than the legitimate winner. Furthermore, there may be collusion between the auctioneer and some of the bidders (similar to shilling).

In general, all electronic schemes in literature can be classified according to how they deal with the trust problem. The following approaches to reducing the trust bidders place in the auctioneer have been identified:

Trusted Third Party

Since the auctioneer is a beneficiary, the assumption that he/she follows the auction protocol may not be realistic. An alternative could be that the bidder and the auctioneer provide a *trusted third party* (TTP) with information so that when there is a dispute the TTP can be called upon to resolve the altercation. However, such a setup requires all parties to have confidence in the TTP. This essentially means that the bidders trade one evil for another. Furthermore, the TTP is an attractive security target and a bottleneck [9].

Threshold Trust

Threshold trust schemes protect against a corrupt auctioneer by distributing the role of the auctioneer across n servers. The auction can be considered secure/fair unless a threshold t , of the auction servers collude (where $t < n$). Threshold trust, however, requires much communication between bidders and the auction servers, as well as between the auction servers themselves [6].

Two-Server Trust

It can be argued that threshold trust is not effective, since collusion among auctioneer servers is beneficial to the whole group of auctioneers. An alternative approach is to split trust up among two servers owned by separate entities.

Here the auction result can be trusted as long as the two entities do not collude. Two-server trust schemes effectively reduce the communication overhead involved in threshold trust schemes and thus far have also proved to be computationally efficient. However, if one of the two servers decide not to co-operate, then the auction outcome cannot be determined. We will give an example of a two-server trust scheme in the next section [7, 13].

Distributed Bidder Trust

In this approach the auctioneer is not used. Instead, the bidders jointly compute the auction outcome. The merit of such an approach is that collusion amongst bidders is prevented unless all bidders are corrupt, which negates the reason for colluding in the first place. However, the downfall of this approach is that **all** bidders must participate during the winner determination stage. This is not reasonable, as when the number of bidders is large (e.g., several hundred) it would be virtually impossible to arrange a time for all bidders to be available to calculate the winner. Furthermore, if just one bidder decides not to participate, the auction outcome cannot be determined. In addition, much communication is required between bidders during the bidding and winner determination stages [1].

Example Electronic Auction

In this section, an English auction scheme based on two-server trust is presented. Note that this example fails some of the security and anonymity criteria stated earlier. It merely serves as an example to highlight the main issues associated with constructing a secure and anonymous electronic auction scheme.

In this example, there are three main parties involved: the bidders, two registration managers (RM1, RM2), and an auctioneer. The auctioneer controls a public bulletin board to which it writes the auction results. All parties can verify the auction proceedings via the bulletin board. Bulletin boards are used in many schemes proposed in literature (see [9, 13]). It is assumed that auction participants use a public key cryptosystem.

Initialization: The auctioneer sets up the auction. The auctioneer and the two registration managers publish their public keys.

Registration: Figure 1 illustrates the process involved in registration.

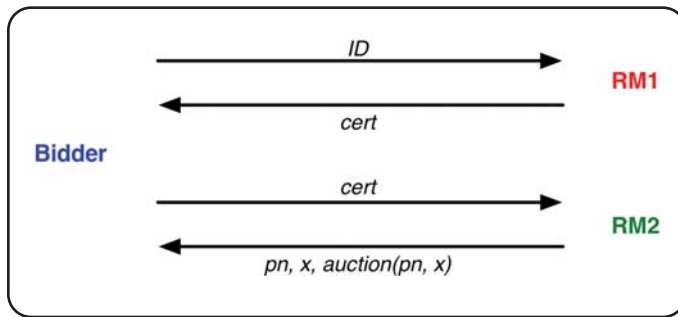


Figure 1: Registration.

1. A bidder registers his/her identity ID with RM1
2. RM1 provides the bidder with a certificate $cert$ (signed using RM1's private key)
3. The bidder presents $cert$ to RM2
4. RM2 checks RM1's signature on $cert$ (using RM1's public key)
5. RM2 issues the bidder with a pseudonym pn , a secret key x , and an auction certificate $auction(pn, x)$ as proof that pn and x are valid

Bidding: Figure 2 illustrates the process involved in bidding.




Figure 2: Bidding.

1. A bidder chooses his/her bid amount $\$$
2. A bidder signs his/her bid $x(\$, pn)$ using x
3. A bidder submits his/her bid to the auctioneer in the following form:

$[\$, pn, x(\$, pn), auction(pn, x)]$

4. The auctioneer checks $auction(pn, x)$ (using RM2's public key)
5. if valid, the bid information is posted on the bulletin board, otherwise, the bid is discarded

Winner Determination: The auctioneer determines the auction outcome according to the auction rules and posts the result to the bulletin board.

Traceability: In the event of a bidder repudiating a bid, an optional traceability procedure can be run:

1. The auctioneer sends the disputed bid to RM1 and RM2
2. RM1 and RM2 recombine their information to reveal the ID corresponding to pseudonym pn that was used to submit the bid. This can be verified by checking $x(\$, pn)$ and $auction(pn, x)$

In this scheme bids are un-forgeable, as they are signed using a bidder's secret key x . Neither of the registration managers is able to frame an individual bidder as they do not know the correspondence between ID , pn , or x . A traceability protocol can be run in the event that a bidder repudiates a bid. The auction proceedings are publicly verifiable via the bulletin board. The scheme is robust in the sense that the auctioneer can discard any bids without valid signatures. Bidders remain anonymous as long as RM1 and RM2 do not collude.

The problem with this example, however, is that bids are linkable because the same pseudonym is used for every bid. While this does not reveal the identity of the bidder, it does over time create a profile about the bidder's bidding strategy. This information could also be used as the basis for a more sophisticated attack (explained in [12]) that allows RM1 to learn the identity of a newly registered bidder. A further problem with this scheme is that RM2 could introduce fake bids into the auction in an attempt to disrupt the auction proceedings with a denial of service attack. Given that RM2 knows the value of every bidder's secret key x , RM2 can randomly sign fake bids using any x .

This simple example showed some of the issues and pitfalls of designing electronic auction schemes.

Research Directions—Group Signatures

The goal of our research is to investigate how group signature schemes can be applied to electronic auctions. The

concept of group signatures was introduced by Chaum and van Heyst [3]. A group signature scheme allows members of a group to sign messages on the group's behalf such that the resulting signature does not reveal their identity. Signatures can be verified with respect to a single group public key, but does not reveal the identity of the signer. Only a designated group manager is able open signatures (reveal the signer's identity), in the case of a later dispute. Furthermore, it is not possible to decide whether two signatures have been issued by the same group member.

Group signatures have been employed by the electronic auction schemes presented in [8, 9]. However, these schemes are very slow, that is, they require many modular computations to be performed and have much communication overhead. Furthermore, the use of group signatures by these schemes is complicated and somewhat different to how more conventional group signature schemes work. Our research focuses on using group signatures in a simpler manner than what has been done in previous schemes. We are also applying this concept to a form of auction referred to as a continuous double auction, where there are many buyers and many sellers (e.g., a share market) [11].

In terms of an auction, the group manager corresponds to a registration manager. The bidders correspond to the group. Once registered, each bidder is able to sign messages (bids) on behalf of the group in such a way that it does not reveal the identity of the individual bidder that submitted the bid. Signatures can be verified by anyone, that is, the auctioneer, other bidders, and outside parties. In the case of a dispute, the registration manager can reveal the identity of the signer of a bid. The role of the auctioneer is merely to organize and run the auction.

Such an approach clearly requires the bidders to have full confidence in the registration manager who is a TTP. However, we are interested in actually distributing the role of the group manager across several parties (as in threshold trust schemes). To the best of our knowledge, group signatures have not been employed in such a manner before. We believe that group signature schemes naturally lend themselves to the electronic auction problem.

Conclusion

Online auctions are inherently insecure. Existing commercial auction sites only offer basic solutions to security and anonymity problems. All individuals are required to trust the auctioneer. Electronic auctions are cryptographic

schemes designed to securely and anonymously conduct auctions. The main goals for a secure and anonymous electronic auction scheme are un-forgable bids (to prevent framing), non-repudiation of bids, public verification of the auction proceedings, robustness, and to protect a bidder's anonymity. It is also highly desirable to reduce the trust that bidders have in the auctioneer. We identified four main approaches to trusting the auctioneer: trusted third party, threshold trust, two-server trust, and distributed bidder trust. We also presented an example of an electronic auction scheme based on the two-server trust approach. This example showed that it is hard to create an electronic auction scheme that satisfies the stated security and anonymity goals. We believe that the good approach to constructing electronic auctions is to use a group signature scheme.

Acknowledgments

The author would like to thank the following people for their contribution to this article: Dr Hossein Ghodosi, Marianne Brown, Bev Trevathan, and Mary Auker.

References

- 1 Brandt, F. Fully private auctions in a constant number of rounds. In the *Seventh Annual Proceedings of Financial Cryptography (FC'03)* (2003).
- 2 Boyd, C., and Mao, W. Security issues for electronic auctions. Hewlett Packard Tech. Rep. TR-HPL-2000 (2000).
- 3 Chaum, D., and van Heyst, E. Group signatures. *Eurocrypt'91* (1991), 257-265.
- 4 Cassady, R. *Auctions and Auctioneering*. University of California Press, Berkeley CA (1967).
- 5 eBay: <http://www.ebay.com>.
- 6 Franklin, M., and Reiter, M. The design and implementation of a secure auction service. *IEEE Trans. Softw. Engin.* 22 (1996), 302-312.
- 7 Naor, M., Pinkas, B., and Sumner, R. Privacy preserving auctions and mechanism design. In the *1st ACM Conference on Electronic Commerce* (1999).
- 8 Nguyen, K., and Traore, J. An online public auction protocol protecting bidder privacy. In the *5th Australasian Conference on Information Security and Privacy (ACISP'00)* (2000), 427-442.
- 9 Omote, K., and Miyaji, A. A practical english auction with one-time registration. In the *6th Australasian Conference on Information Security and Privacy (ACISP'01)* (2001), 221-234.

- 10 onSale: <http://www.onsale.com>.
- 11 Trevathan, J. An anonymous and secure continuous double auction scheme. School of Information Technology, James Cook University Tech. Rep. (2004).
- 12 Trevathan, J., and Ghodosi, H. Design issues for electronic auctions. School of Information Technology, James Cook University Tech. Rep. (2004).
- 13 Wang, C., and Leung, H.-F. Anonymity and security in continuous double auctions for Internet retailers market. In the *37th Hawaii International Conference on System Sciences* (2004).
- 14 Vickrey, W. Counter speculation, auctions and competitive sealed tenders. *J. Finance* 16 (1961), 8-37.

Biography

Jarrold Trevathan (jarrod@cs.jcu.edu.au) is a PhD student at the School of Information Technology, James Cook University. His research interests include cryptography, database security, multimedia, and electronic commerce.

Call for Articles

CROSSROADS

THE ACM STUDENT MAGAZINE

ISSUE 12.2 • WINTER

“Human Computer Interaction”

Due 1 August 2005

ISSUE 12.3 • SPRING

“Programming Languages”

Due 4 October 2005

www.acm.org/crossroads/doc/cfas

Computing is the future. Is it yours?

Earn your M.S. or Ph.D. without interrupting your career.

We have a beautiful, 300-acre main campus, but that's only the beginning. We also have a magnificent 57,800,000-square-mile online campus that really works for working professionals.

- Leaders in online education since 1983
- Listed in The Princeton Review's, *The Best Distance Learning Graduate Schools*
- “One of the nation's top 20 cyber-universities”
—*Forbes Magazine*
- Information security curriculum certified by the U.S. National Security Agency (NSA)
- Upsilon Pi Epsilon Honor Society
- Online or on-campus

Graduate School of Computer and Information Sciences

www.scis.nova.edu
scisinfo@nsu.nova.edu
800-986-2247, ext. 2000



Computer
Information
Systems

Computer
Science

Computing
Technology
in Education

Information
Science

Information
Security

Information
Systems

3.2 Design Issues for Electronic Auctions

This paper investigates cryptographic protocols for privacy and security in auctions. Existing proposals for auction security are reviewed, and the differing security requirements between various auction types are contrasted. The paper illustrates design flaws in several existing ‘secure’ auction schemes. This paper introduces the main theme behind the thesis’ overall security solution for English auctions and CDAs. The proposed solution is based on a *group signature scheme*. Incorporating a group signature scheme into an auction allows for anonymous and verifiable bid submission. A bidder’s identity can be recovered in the situation where a bid is repudiated. The protocol details specific to English auctions and CDAs are described in Chapters 4 and 7 respectively.

“*Design Issues for Electronic Auctions*” [4] is published in the 2nd International Conference on e-Business and Telecommunications Networks 2005 (ICETE). ICETE 2005 was held in Reading, United Kingdom and was sponsored by MICROSOFT UK and Reading University. The major goal of this conference was to bring together researchers and developers from academia and industry working in areas related to e-business, to stimulate a debate with a special focus on telecommunications networks. The conference received 151 papers in total, with contributions from 38 different countries. To evaluate a submission, a double blind paper evaluation method was used: each paper was reviewed by at least two internationally known experts from the Program Committee. In the end, 85 papers were selected for oral presentation and for publication, corresponding to an acceptance rate of 56%. This paper is available online via citeseer⁵, and referenced by DBLP⁶.

I was a session chair at ICETE 2005. This paper is co-authored by Dr Hossein Ghodosi. Dr Ghodosi’s research interests include Cryptography and Information Security. Publication was made possible by the Graduate Industry and Travel Awards (GRITA) and the School of Mathematical and Physical Sciences.

⁵<http://citeseer.ist.psu.edu/>

⁶<http://www.informatik.uni-trier.de/~ley/db/>

DESIGN ISSUES FOR ELECTRONIC AUCTIONS

Jarrold Trevathan and Wayne Read

School of Mathematical and Physical Sciences

James Cook University

Email: jarrod.trevathan/wayne.read@jcu.edu.au

Hossein Ghodosi

School of Information Technology

James Cook University

Email: hossein@cs.jcu.edu.au

Keywords: Auction security, anonymity, authentication, price determination, payment enforcement, group signatures

Abstract: Extensive research has been conducted in order to improve the security and efficiency of electronic auctions. However, little attention has been paid to the design issues. This paper discusses design issues and contrasts the differing security requirements between various auction types. We demonstrate that poor design for an electronic auction breaches the security of the system and degrades its practicality, irrespective of how secure/efficient the building blocks of an electronic auction are. This is accomplished by illustrating design flaws in several existing electronic auction schemes. Furthermore, we provide a solution to these flaws using a group signature scheme and give recommendations for sound auction design.

1 INTRODUCTION

An auction is an exchange mechanism whereby a seller (or sellers) offers an item for sale and many bidders submit bids in competition for the item. The Auctioneer organises the auction on behalf of the seller and accepts bids from the bidders. The Auctioneer attempts to maximise the sale price for the seller, whereas the bidders try to win the item for the lowest price possible. An electronic auction is a cryptographic protocol designed to securely and anonymously conduct auctions. While extensive research has been undertaken for improving the security and efficiency of electronic auctions, (see (Cachin, 1999; Franklin and Reiter, 1996; Kikuchi *et al.*, 1998; Naor *et al.*, 1999; Trevathan, 2005)), little attention has been paid to the design issues. In this paper we demonstrate that poor design for an electronic auction breaches the security of the system and degrades its practicality regardless of how secure/efficient the building blocks of an electronic auction are. This is accomplished by illustrating design flaws in several existing auction schemes. We also provide solutions to these flaws using a group signature scheme and give some recommendations for sound auction design.

This paper is organised as follows: Section 2 gives a general background on electronic auction schemes. Section 3 describes the security problems inherent in

conducting electronic auctions. Section 4 discusses design issues and illustrates some major design flaws in several existing schemes. Section 5 gives some recommendations about how these flaws can be fixed and suggestions regarding good auction design principles. Section 6 provides some concluding remarks.

2 TYPES OF AUCTIONS

There are many types of auctions, including English, Vickrey, Dutch, and Continuous Double auctions (see Figure 2). An English auction is the most prominently known type of auction and is commonly used in real estate. In this auction there is one seller who offers an item for sale. Many potential buyers submit bids in an attempt to win the item. The winner is the highest bidder after a given time-out period.

An alternative to this is the Vickrey auction. This is referred to as a sealed bid, second price auction. It is comprised of bidders submitting their bids secretly during a bid submission round. The bids remain sealed until a winner determination round, where the Auctioneer views all of the bids and awards the auction to the highest bidder. In this auction, the winner only has to pay the second highest price, i.e., the highest losing bid. The values of all other losing bids remain secret. In contrast, an English auction is an open bid, first price auction.




Figure 1: Taxonomy of Electronic Auction Schemes

In a Dutch or descending auction, the seller progressively lowers his/her bid until a buyer accepts. The winner is the first buyer to accept the seller's offer. They must pay an amount equal to this bid. In this respect Dutch auctions have the natural property of concealing the losing bid values. This form of auctioning is also referred to as descending, as unlike an English auction, the winning price is bid downwards rather than up.

A further style of auction used in share markets is a Continuous Double Auction (CDA). The aforementioned auctions have only one seller and many buyers. CDAs on the other hand, have many buyers and many sellers who continuously trade a particular commodity. These auctions are open bid and the price can be either ascending or descending.

Electronic auction schemes have been proposed for all of these auction types. However, regardless of the auctioning mechanism, all schemes consist of the following stages:

Initialisation: The Auctioneer sets up the auction and advertises it (i.e., description of good, starting time, etc).

Registration: In order to participate in the auction, bidders must first register with the Auctioneer (or a registration manager). This ensures only valid bids are made and bidders can be identified for payment purposes. It is desirable for registration to be a one-off procedure. When a bidder has registered they are able to participate in any number of auctions rather than re-registering for each new auction.

Bidding: A registered bidder computes his/her bid and submits it to the Auctioneer. The Auctioneer checks the bid received to ensure conformity with the auction rules.

Winner Determination: The Auctioneer determines the winner according to the auction rules. It is desirable

for this process to be publicly verifiable.

3 AUCTION SECURITY

In recent years, many companies have emerged offering auctioning services via the Internet (e.g., eBay¹ and onSale²). Such sites lack security for both the seller and the bidder and moreover require all parties to trust the Auctioneer. Furthermore, the identities and information relating to the participants is open to abuse. Common problems identified include: the bidders repudiating bids (i.e., they win and later decide they don't want to pay), the seller not delivering the item and Auctioneer corruption by awarding the auction to someone other than the legitimate winner.

Security and anonymity has been addressed in literature by (Cachin, 1999; Franklin and Reiter, 1996; Kikuchi *et al.*, 1998; Naor *et al.*, 1999; Trevathan, 2005). Additionally, numerous auction schemes have been proposed in an attempt to solve the aforementioned problems. There are differing security requirements depending on whether the auction is sealed or open bid. However, in general, most schemes in literature agree on the following security goals:

Unforgeability - Bids must be unforgeable, otherwise a bidder can be impersonated.

Non-Repudiation - Once a bidder has submitted a bid they must not be able to repudiate having made it. For example, if a bidder wins and does not want to pay they might deny that they submitted the bid.

Anonymity - The bidder-bid relationship must be

¹<http://www.ebay.com>

²<http://www.onsale.com>

concealed so that no bidder can be associated with the bid they submit.

Public Verifiability - There must be publicly available information by which all parties can be verified as having correctly followed the auction protocol. This should include evidence of registration, bidding and proof of winner/loser.

Robustness - The auction process must not be affected by invalid bids nor by participants not following the auction protocol correctly.

4 DESIGN ISSUES

The main design goals discussed previously in the literature of electronic auctions include fairness, efficient registration, efficient bidding, and efficient winner determination (see, for example, (Boyd and Mao, 2000; Franklin and Reiter, 1996; Naor *et al.*, 1999; Boyd *et al.*, 2000; Trevathan, 2005)). Our observation is that despite all components of an electronic auction functioning correctly, there is still no guarantee that the whole system works properly. In this section we discuss several design issues that must be considered in order to achieve a practical auction system.

4.1 Trust Issues

Auctioneer corruption is a major problem in electronic auctions. A malicious Auctioneer might influence the auction proceedings in a manner inconsistent with the auction rules. For example, the Auctioneer might choose to block bids, insert fake bids, steal payments, profile bidders, open sealed bids prior to the winner determination phase, or award the item to someone other than the legitimate winner. Furthermore, the Auctioneer may be in collusion with some of the bidders.

In general, all schemes can be classified according to how they deal with this problem. We have identified the following trust models:

Auctioneer Trust - This is the easiest (and the most unacceptable) solution to the problem. It is assumed that the Auctioneer is trustworthy and honestly follows the protocol. This is the strategy employed by all Internet auction sites as it easily solves many relevant problems (e.g., security, anonymity, etc.). However, all bidders are at the mercy of the company running the auction.

Trusted Third Party - Since the Auctioneer is a beneficiary, the assumption that it follows the auction protocol is unrealistic. An alternative could be that the bidders and Auctioneer provide a *trusted third party* (TTP) with information so that when there is a

dispute the TTP can be called upon to resolve it (see (Boyd *et al.*, 2000)). However, such a setup requires all parties to have confidence in the TTP who is an attractive security target and a bottleneck.

Threshold Trust - Threshold trust schemes protect against a corrupt Auctioneer by distributing the role of the Auctioneer across ℓ servers (see (Franklin and Reiter, 1996; Sakurai and Miyazaki, 1999)). The auction can be considered secure/fair unless a threshold t , $1 \leq t \leq \ell$ of the Auctioneers collude. The value of t is usually around $\ell/3$.

Such a scheme is clearly better than the previous approaches as no single party acting maliciously can influence the auction proceedings. However, threshold trust requires much communication between bidders and the auction servers, as well as between the auction servers themselves. Furthermore, when the auctioning company is small, collusion among Auctioneers is beneficial to the group as a whole.

Two-Server Trust - An alternative approach to threshold trust is to split the auctioning responsibility among two servers owned by separate entities (see (Cachin, 1999; Naor *et al.*, 1999)). Here the auction result can be trusted as long as the two entities do not collude. Two-server trust schemes effectively reduce the communication overhead involved in threshold trust schemes and thus far have proven to be computationally efficient. However, if one of the two servers decides not to co-operate, then the auction outcome cannot be determined.

Distributed Trust - In this approach the bidders jointly calculate the auction result without the help of an Auctioneer. The merit of such an approach is that collusion amongst bidders is prevented unless all bidders are corrupt, which negates the reason for colluding in the first place (see (Brandt, 2003)).

Distributed trust schemes are unrealistic as they require all bidders to participate during the winner determination phase. Such an approach is not feasible when the number of bidders is large. Furthermore, distributed trust is not applicable to CDAs as this assumes knowledge of the total number of bidders prior to the start of bidding. In a CDA, new bidders are entering the auction all the time. Using a distributed bidder approach would require communicating new information to all the auction participants every time a new bidder registers.

4.2 Anonymity Issues

Another important criteria for the evaluation of secure electronic protocols is anonymity. This is important because the bidder-bid relationship must be concealed in such a way that no bidder can be associated with the

bid they have submitted. There are several different approaches to satisfy this requirement. For example, auction protocols which utilise secure computation for winner determination (e.g., (Kikuchi *et al.*, 1998; Harkavy *et al.*, 1998) provide this facility by concealing the identity of bidders in their bids). A common solution is to issue bidders with a pseudonym (during registration), which they can use to submit bids. That is, bidders register themselves (by presenting verified identification) and obtain a pseudonym.

Seemingly the technique of issuing pseudonyms requires having more than one server/party on the auction side. Otherwise, the same party issuing the pseudonym knows the real identity of the associated bidder, and thus learns the relationship between the bids and the bidders. On the contrary, if the pseudonym issuer does not know the pseudonym (e.g., it is issued by blindly signing a message), then it cannot retrieve the real identity of a bidder in the case of a dispute. This model essentially works as follows (assume the two servers/parties are called S_1 and S_2):

1. Bidders present their identification to S_1 and obtain a token that does not carry their ID.
2. Bidders submit the token (without revealing their ID) to S_2 , who issues a pseudonym associated with the token.

In this way, neither S_1 nor S_2 knows the relationship between any real ID and the pseudonyms. The bidder then submits his/her bid using the pseudonym. However, in the event of a dispute S_1 and S_2 can cooperate to determine the real ID associated with each pseudonym.

Our observation is that, regardless of how secure the anonymity issuing protocol is, the resulting scheme is not secure if there is no separation between the registration and the bidding phases. That is, the registration must be performed for all bidders prior to the commencement of bidding and the system must not accept any bid before the registration is closed. The following scenario explains a possible attack: A bidder provides identification to S_1 and obtains a token. Using this token, it obtains a pseudonym which can be used for bidding. If there is no separation between the registration and bidding phases S_1 can act in a procrastinating manner by halting all future registrations until the newly registered bidder submits his/her bid. This scenario enables S_1 to learn the mapping between the bidder's identity and his/her pseudonym. To protect against this type of attack, the scheme should not allow any bidding prior to the registration closing time.

Note that there are electronic auction schemes in which it is impossible to have a separation between the registration and the bidding phases. For example, a CDA allows bidders to continuously submit bids at the same time new bidders are being registered (i.e.,

the registration and bidding phases overlap). Therefore the scheme by Wang and Leung (Wang and Leung, 2004) can be broken using this *procrastinating attack*.

4.3 Bid Authentication Issues

Efficient winner determination is an important criteria in the evaluation of electronic auction protocols. Because of its importance, many schemes provide evidence to support their claims regarding the ability of their system to efficiently process bids. For example, (Franklin and Reiter, 1996) in their highly referenced work, claim that:

“We have implemented a prototype of our service to demonstrate its feasibility. The performance of this implementation indicates that our approach is feasible using off-the-shelf workstations for auction servers, even for large auctions involving hundreds of bids.”

The question is, what will happen if a (set of) malicious bidder(s) issues too many bids? Optimistically assuming that processing each bid takes only one second, then the winner determination process will require a proportional amount of time (i.e., the system is not practical). The problem is more crucial in schemes which use secure computation (see, e.g., (Kikuchi *et al.*, 1998)). They achieve anonymity by concealing the identity of bidders in their bids. In addition, all bids are submitted anonymously. The winner determination protocol opens only one bid – a bid that contains the highest offer. If there is a tie (i.e., more than one bid at the highest offer), then another round of bidding must occur. A malicious bidder can easily cause a *never ending scenario* in this scheme. For example, even if all high valued bids are opened (which is disallowed by the protocol), the Auctioneer cannot determine who has submitted the bid as it contains a false identity. Obviously, such schemes are not practical at all.

This problem can be avoided if the bids have been authenticated. That is, the Auctioneer accepts bids only from registered bidders (generally, in all sealed-bid auctions, each bidder submits only one bid). If the system supports anonymity of the bidders, then it must also provide the authentication of the corresponding pseudonym. Note that there are schemes which check the validity of the bids (i.e., they check whether the submitted bid satisfies a predetermined structure). This is insufficient, as one may submit too many well-structured bids. Furthermore, bid authentication must be secure about relevant attacks. In order to illustrate the problem, let us examine the protocol by (Boyd *et al.*, 2000), which supports a sealed bid auction system. This is possibly the only auction scheme which uses bid

Bidder	Auctioneer
$a, d_1, d_2 \in_R Z_q^*$ $S = g_1^b g_2^a, B = g_1^{d_1} g_2^{d_2}$	$\xrightarrow{S, B}$ \xleftarrow{c}
$s_1 = d_1 - cx_1, s_2 = d_2 - cx_2$ $t_1 = s_1 - bc, t_2 = s_2 - ac$	$c \in_R Z_q$ $\xrightarrow{t_1, t_2}$
	$B \stackrel{?}{=} (S y_1 y_2)^c g_1^{t_1} g_2^{t_2}$

Table 1: The Sealing Protocol

authentication. But, as we will show shortly, it is subject to an *impersonation attack*.

System Settings in (Boyd et al., 2000) -

G is a subgroup of order q in Z_p^* , such that $p = 2q + 1$ for sufficiently large prime p . There are two generators, g_1 , and g_2 , such that nobody knows $\log_{g_1} g_2$. The public keys of a bidder b_i are certified to be $y_1 = g_1^{x_1}$ and $y_2 = g_2^{x_2}$.

Sealing Protocol in (Boyd et al., 2000) -

This is an interactive protocol between the bidder and the Auctioneer. At the end of this protocol, the Auctioneer will be convinced that the bidder submitted a correct bid which can be opened at the bid-opening phase (in their scheme, the help of the bidder is necessary to open the bid). A bidder, b_1 , with public keys $y_1 = g_1^{x_1}$ and $y_2 = g_2^{x_2}$ wishes to commit himself/herself to the bid value b . The protocol has three steps and works as shown in Table 1.

To open the bid (at the bid-opening phase) the bidder should release the tuples (b, a) . This tuple can be checked with corresponding values at the bid sealing protocol, and upon verification, the value b is the bid submitted by bidder b_1 . If a bidder refrains from opening his/her bid (e.g., when he/she learns that his/her bid is too much higher than the others), the identity of this person can be traced and suitable action will be taken. The idea behind this protocol is that the tuple (S, B, c, t_1, t_2) is unique, and thus the bidder is committed to the bid value b . The authors also proposed a non-interactive version of this protocol (see (Boyd et al., 2000)).

Anyone that does not know the secret keys x_1 and x_2 , cannot participate in, and complete the protocol successfully. The impersonation attack works as shown in Table 2 (the non-interactive version is also subject to this attack).

The consequence of this attack is that an innocent bidder b_1 will be accused of not participating in the bid opening protocol. But, in fact, nobody knows the relevant tuple (b, a) for this case (if solving the Dis-

Bidder	Auctioneer
$t_1, t_2 \in_R Z_q^*$ $S = \frac{1}{y_1 y_2}, B = g_1^{t_1} g_2^{t_2}$	$\xrightarrow{S, B}$ \xleftarrow{c} $\xrightarrow{t_1, t_2}$
	$c \in_R Z_q$ $B \stackrel{?}{=} (S y_1 y_2)^c g_1^{t_1} g_2^{t_2}$

Table 2: The Sealing Protocol by a Malicious Bidder

crete Logarithm problem is hard).

4.4 Price Determination Issues

Many auction schemes impose a price list on bidders (e.g., (Harkavy et al., 1998; Sakurai and Miyazaki, 1999)). This is essentially a series of bidding points (w_1, w_2, \dots, w_k) where $w_1 < w_2 < \dots < w_k$ which a bidder must choose from.

The scheme by (Harkavy et al., 1998) requires bidders to submit a bid value for each point in the price list. Meaning, for each point in the price list, a bidder must either bid his/her *id* or a nullifying value. For a price list of size k , this requires a bidder to submit k individual values. In some types of auctions (such as CDAs), a price list may be too restrictive as bidders require flexibility in choosing their bid. In some instances this may be down to fractions of cents. Adding more bidding points comes at the cost of increased computation and communication overhead. So the question here is what is the correct value of k to choose?

Moreover, two colluding bidders can continually bid at the highest price level and cause infinite rounds of auctioning (i.e., *denial of service attack*). This is a major problem with the schemes of (Harkavy et al., 1998) and (Sakurai and Miyazaki, 1999). For example, if the highest price in the list is k and two bidders bid at this price, then the auction goes to a second round of bidding. In this round the maximum value in the price list increases to $k \times 2$. However, both bidders then bid $k \times 2$. This then forces a third round of bidding and the maximum value in the price list increases to $k \times 3$, etc. If the bidders continue in this manner, the auction continues indefinitely.

In addition, many schemes using a price list, also require bidders to participate in a disavowal protocol (e.g., (Sakurai and Miyazaki, 1999)). This requires the bidder to engage in a protocol with the Auctioneer to show that the bid they submitted is not for that particular price. Upon completion of disavowal for the respective price level, the Auctioneer selects the next price in the list and the process repeats. Unfortunately disavowal protocols are very expensive as much communication must take place to determine the winner.

Furthermore, if some bidders do not participate they are instantly incriminated, or the auction process is left in an inconsistent state.

For the reasons stated above, schemes using price lists tend to be restrictive and are not always practical for certain types of auctions.

4.5 Payment Issues

A major issue in electronic auction schemes is how to enforce payment from the winning bidder(s). Previous schemes have achieved this by using digital cash schemes. For example, Franklin and Reiter (Franklin and Reiter, 1996) use a digital coin to represent a bidder's bid. The value of the coin is equivalent to the amount of the bid. The winner's coin is deposited into the seller's bank account at the end of the protocol. In this scheme the purpose of the coin is to prevent repudiation of bids. The scheme by (Boyd *et al.*, 2000), uses a digital cash scheme in a similar manner. When a bidder repudiates a bid, the identity of the bidder can be recovered by presenting the bidder's piece of digital cash to the bank.

We believe that using digital cash does not significantly enhance the security of an electronic auction scheme. Instead it introduces another party (i.e., the bank) to the auction which must be trusted by the participants. This complicates auction schemes and makes it increasingly difficult to analyse how secure the scheme is. Franklin and Reiter (Franklin and Reiter, 1996) claim that their scheme can provide anonymity through the use of pseudonyms. However even if this is done, anonymity is still dependent on the underlying digital cash scheme.

The main goal of digital cash is to allow a spender to engage in anonymous and untraceable transactions with a merchant. When the spender breaks the rules (i.e., repudiates the purchase or double spends the cash) the only means of recourse is for the bank to reveal the spender's identity. In terms of an anonymous auction, we are only interested in revealing a bidder's identity when there is a dispute (i.e., bid repudiation). Unless the bank is built into the auctioning model (as in (Boyd *et al.*, 2000)), other more efficient mechanisms for identity escrow can be used rather than digital cash schemes.

5 RECOMMENDATIONS

This section provides some suggestions of ways to fix the problems raised in the previous section and discusses good design strategies for electronic auctions. Surprisingly there already exists cryptographic mechanisms to solve many of the problems facing

electronic auctions. However, little if any of the literature on electronic auctions has embraced this research. The proposed solution comes from group signature schemes.

A group signature scheme allows members of a group to sign messages on behalf of the group so that no one can work out which particular user is the signer (see (Anteniese *et al.*, 2000)). A signature on a message can be verified by anyone using a single public key. When there is a dispute (i.e., someone repudiates having signed a message) a trusted group manager can reveal the identity of the signer of a message. Each signature is unlinkable and no coalition of users (even with the help of the group manager) is able to forge a valid signature of an innocent user.

Group signatures have many desirable properties with regard to electronic auctions. In an auction, each bidder belongs to a group of bidders. Every bidder has a unique signature which enables them to sign bids anonymously. The bidders submit their bids to the Auctioneer, who can then verify each signature using the group's public key. However, the Auctioneer cannot solely be the group manager as this would require all bidders to implicitly trust that the Auctioneer would not reveal their identities. Instead this task must be distributed among the Auctioneer and another party (e.g., a registration manager). When there is a dispute, both the Auctioneer and the other party must cooperate in order to reveal the identity of whom signed the bid in question.

In the subsections that follow, we will address each of the design issues stated in the previous section and describe how the properties of group signature schemes can be employed. Throughout this section, we develop a model for electronic auction design. This model is not specific to any particular auction type, but can be used as a general basis for all types of auctions. Note that mechanisms for sealing bids will not be discussed. Instead the model addresses trust, anonymity, bid authentication, price flexibility and payment enforcement.

5.1 Trust

Using publicly verifiable protocols reduces the trust required in the Auctioneer. This approach allows all parties (even outsiders) to view the auction process and they can verify that the auction outcome is correct.

Many schemes use public bulletin boards to serve this purpose (Sakurai and Miyazaki, 1999; Boyd *et al.*, 2000; Wang and Leung, 2004). An Auctioneer can post bids and auction results on the board so that all others can verify the auction proceedings. It is assumed that only the Auctioneer can write to this board and that it is publicly accessible to all parties.

Public bulletin boards have implications for sealed bid auctions as the main requirement is to keep the losing bids secret. However, in all sealed bid auction schemes it is inevitable that bidders learn statistics about the bids submitted (i.e., highest price, bid values, etc). Therefore, using a bulletin board is sufficient for sealed bid auctions, as long as the bidder-bid relationship remains secret (i.e., bidders cannot be linked to their bids).

However, even if publicly verifiable protocols are used there is still some need to trust the Auctioneer. In this case, the best approach is to use the two-server trust model (see Section 4.1). This needs to be modified slightly in that the second party is only required in the case of a dispute but cannot act on his/her own (as in the case of the TTP approach). This will be clarified in the next section.

5.2 Anonymity

To protect against the procrastinating attack described in section 4.2, there must be complete separation between the registration and bidding stages of an auction. This means that all bidders must register prior to the Auctioneer accepting any bids. This is a straightforward matter for English, Vickrey and Dutch auctions as the Auctioneer can prevent the auction from starting until all bidders have registered.

However, in CDAs it is impossible to separate the registration and bidding stages. In this situation a group signature scheme can be used. Group signatures are unlinkable, which means that given two signatures, it is impossible to tell if the same user signed both messages. This thwarts the procrastinating attack as S_1 can no longer determine if the signature that they are observing belongs to the newly registered bidder or someone else. Even if S_1 does learn the mapping between the newly registered bidder and the signature on a bid, the registration manager will not be able to link future bids to the bidder due to the unlinkability property of the group signature.

5.3 Bid Authentication

In general, authentication protocols require considerable time consuming operations such as exponentiation. However, there are algorithms that can be employed for performing multiple verifications in one stage/operation (e.g., (Bellare *et al.*, 1998)). Group signatures are also an effective method of performing bid authentication.

5.4 Price Determination

Bidders must be allowed flexibility when choosing their bid values. Using a group signature approach fa-

cilitates price flexibility as bidders are free to choose whatever bid values they desire. Furthermore, a group signature scheme does not require the use of a disavowal protocol. In general, designers of electronic auctions should avoid disavowal protocols.

5.5 Payment Enforcement

As stated in Section 4.5, digital cash schemes only serve as a non-repudiation mechanism in auctions and makes the scheme conceptually harder to understand. There are more efficient identity escrow mechanisms other than digital cash. Again group signature schemes are an ideal choice. When there is a dispute, (i.e., someone repudiates a bid, or refuses to pay), the group manager can reveal the identity of the person that signed the bid in question. This is more efficient than a digital cash scheme and serves the same purpose.

5.6 Group Signature Auction Model

This section proposes a generic auction model using a group signature scheme. The model incorporates the recommendations from the previous subsections. Furthermore, it gives the designer flexibility to use any group signature scheme.

There are three parties in the auction:

1. A Bidder, who is interested in submitting bids for an item offered by a seller.
2. An Auctioneer, who organises the auction, accepts the bids and determines the winner according to the auction rules. The Auctioneer also holds the corresponding relation of the identity and a token associated with each bidder. The Auctioneer has access to a publicly verifiable bulletin board.
3. A Registrar, who takes part in a protocol in order to complete the registration of a bidder who has obtained a token from the Auctioneer. At the end of the protocol, the bidder obtains a secret key that enables him/her to generate signed bids in a proper format.

Setup - The Auctioneer organises the auction (i.e., advertising and calls for auction). The Registrar sets up the group public key and his secret key.

Registration - A user submits a request to the Auctioneer to participate in the auction. The Auctioneer verifies the identity of the requestor, and issues a token that is verifiable by the Registrar. The user then takes part in a protocol with the Registrar, in order to obtain his secret key and a certificate of membership in the auction. Note that the token does not carry the real identity of the bidder, but all communication between the Registrar and the owner of a token is authenticated, and will be kept for tracing, or revealing the identity of users associated with tokens.

Bidding - Using a membership certificate, a bidder can generate anonymous and unlinkable group signatures on a bid. A bidder submits a bid to the Auctioneer signed using the secret key. The Auctioneer verifies the signature on the bid using the public key. If the bid is valid, the Auctioneer posts it on the bulletin board.

Winner Determination - The Auctioneer determines the auction outcome according to auction rules (i.e., English, Vickrey, etc.). The auction result is posted on the bulletin and the winner can produce his/her signed bid as evidence that they won.

Traceability - In the event of a dispute (i.e., bid repudiation), the Auctioneer and the Registrar can combine their information to reveal the identity of a signer.

6 CONCLUSION

Limited attention has been paid to the design issues of electronic auctioning schemes. This paper demonstrated that poor auction designs breach the security of any auction scheme irrespective of how secure the underlying cryptographic building blocks are. Auctions were examined in terms of trust, anonymity, bid authentication, price determination and payment enforcement.

We have shown that many schemes proposed in literature are not practical and can be broken by exploiting weakness in the fundamental design employed. A group signature scheme can solve many of the problems in electronic auction design. However, this is not an all-encompassing solution to all of the design issues. It is insufficient to rely on cryptographic mechanisms alone. An auction system is only secure as its weakest component.

There still remain many pressing issues in designing electronic auctions. Now the time has come for the designers of auction schemes to step away from the small-view approach and observe the auction system as a whole. Designers need to analyse how each component interacts and scrutinise whether the cryptographic methods employed are really sufficient for use in electronic auction schemes.

REFERENCES

- G. Ateniese, J. Camenisch, M. Joye, and G. Tsudik, "A Practical and Provably Secure Coalition-Resistant Group Signature Scheme," in *Advances in Cryptology - Proceedings of CRYPTO 2000* (M. Bellare, ed.), vol. 1880 of *Lecture Notes in Computer Science*, pp. 255–270, Springer-Verlag, 2000.
- M. Bellare, J. Garay, and T. Rabin, "Fast Batch Verification for Modular Exponentiation and Digital Signatures," in *Advances in Cryptology - Proceedings of EUROCRYPT '98* (K. Nyberg, ed.), vol. 1403 of *Lecture Notes in Computer Science*, pp. 236–250, Springer-Verlag, 1998.
- C. Boyd and W. Mao, "Security Issues for Electronic Auctions," tech. rep., Hewlett Packard, TR-HPL-2000, 2000.
- F. Brandt, "Fully private auctions in a constant number of rounds," in *Proceedings of the 7th Annual Conference on Financial Cryptography (FC)* (R. Wright, ed.), vol. 2742 of *Lecture Notes in Computer Science*, pp. 223–238, Springer-Verlag, 2003.
- C. Cachin, "Efficient Private Bidding and Auctions with an Oblivious Third Party," in *6th ACM Conference on Computer and Communication Security*, pp. 120–127, 1999.
- M. Franklin and M. Reiter, "The Design and Implementation of a Secure Auction Service," *IEEE Transactions on Software Engineering*, vol. 22, pp. 302–312, May 1996.
- M. Harkavy, J. Tygar, and H. Kikuchi, "Electronic Auctions with Private Bids," in *the 3rd USENIX Workshop on Electronic Commerce*, Aug. 1998.
- H. Kikuchi, M. Harkavy, J. Tygar, "Multi-round Anonymous Auction Protocols," *IEEE Workshop on Dependable and Real-Time E-Commerce Systems*, pp. 62–69, 1998.
- M. Naor, B. Pinkas, and R. Sumner, "Privacy Preserving Auctions and Mechanism Design," in *The 1st Conference on Electronic Commerce*, pp. 129–139, 1999.
- K. Sakurai and S. Miyazaki, "A Bulletin-Board Based Digital Auction Scheme with Bidding Down Strategy," in *International Workshop on Cryptographic Techniques and E-Commerce*, pages 180–187, 1999.
- J. Trevathan, "Security, Anonymity and Trust in Electronic Auctions," *Association for Computing Machinery, Crossroads Magazine*, Spring Edition, vol. 11.3, 2005.
- K. Viswanathan, C. Boyd, and E. Dawson, "A Three Phased Schema for Sealed Bid Auction System Design," in *Proceedings of ACISP 2000 - Australasian Conference on Information Security and Privacy* (E. Dawson, A. Clark, and C. Boyd, eds.), vol. 1841 of *Lecture Notes in Computer Science*, pp. 412–426, Springer-Verlag (Berlin), 2000.
- C. Wang and H. Leung, "Anonymity and Security in Continuous Double Auctions for Internet Retail Market," in *the 37th Hawaii International Conference on Systems Sciences*, 2004.

Chapter 4

Privacy and Security in English Auctions

English auctions are the most common form of online auction. Previous cryptographic solutions have mainly concentrated on other auction types such as Sealed Bid and Dutch auctions. While these schemes are helpful in defining the basic auction security requirements, they are of limited usefulness to online English auctions. This chapter presents one paper that proposes a cryptographic scheme that is tailored to online English auctions (see Section 4.1). CDA security is addressed in Chapter 7.

4.1 Secure Online English Auctions

Online auctions such as those used by eBay are based on a type of auction referred to as an English auction. This paper proposes a new scheme for conducting secure and anonymous online English auctions using a modified type of group signature. Trust is divided among three servers owned by separate companies to ensure anonymity and fairness. The proposed scheme solves the problems of the existing English auction schemes and has following characteristics: *unforgeability*, *anonymity*, *unlinkability*, *exculpability*, *coalition-resistance*, *verifiability*, *robustness*, *traceability*, *revocation* and *one-off registration*. The scheme can prevent the Auctioneer from altering the auction timing (*unskewability*) and from selectively excluding bids (*unblockability*). The proposed scheme has comparable efficiency to the existing schemes for the enhanced privacy and security it provides.

“*Secure Online English Auctions*” [7] is published in the International Conference on Security and Cryptography Conference 2006 (SECRYPT). SECRYPT 2006 was held in Setúbal, Portugal and was organised by INSTICC. Sponsors included IBM, Setúbal Polytechnic Institute, IEEE Systems, Man and Cybernetics (SMC) Society, and the International Association for Cryptographic Research (IACR). The major goal of this conference was to bring together researchers and developers from academia and industry working in areas related to e-business to stimulate a debate with a special focus on telecommunications networks. The conference received 326 submissions in total, with contributions from 53 different countries. To evaluate each submission, a double blind paper evaluation method was used: each paper was reviewed by at least two internationally known experts from the Program Committee, and more than 95% of the papers had 3 or more reviews. In the end, 98 papers were selected to be published and presented

as full papers. This corresponds to a 30% acceptance ratio. This paper is available online via citeseer ¹ and is referenced by DBLP ².

Publication in SECRYPT 2006 was made possible by funding from the School of Mathematical and Physical Sciences.

“*Secure Online English Auctions*” [18] has also been selected to appear in the “ICETE 2006 Best Papers” book published by Springer Verlag. SECRYPT is part of the International Conference on E-Business and Telecommunications Engineers (ICETE). Out of four independent conferences held under the ICETE banner, only 28 of the best papers were accepted for publication in the book. This corresponds to an 8% acceptance rate.

¹<http://citeseer.ist.psu.edu/>

²<http://www.informatik.uni-trier.de/~ley/db/>

SECURE ONLINE ENGLISH AUCTIONS

Jarrold Trevathan

School of Mathematical and Physical Sciences

James Cook University

Email: jarrod.trevathan@jcu.edu.au

Wayne Read

School of Mathematical and Physical Sciences

James Cook University

Email: wayne.read@jcu.edu.au

Keywords: Online auctions, event timing, anonymity, group signature.

Abstract: Security and privacy in online auctions is a major concern as auction participants have many opportunities to cheat (e.g., repudiate bids, not deliver items, etc.). Online auctions such as those used by eBay are based on a type of auction referred to as an English auction. Despite the English auction being the most popular type of auction, it has received less security coverage than other types of auctions (e.g., sealed-bid auctions). An existing proposal for a “secure” English auction prevents the Auctioneer from closing the auction early and from blocking bids, but does not protect a bidder’s anonymity. Another proposal provides anonymity, but does not stop an Auctioneer from skewing its clock or blocking bids. This paper proposes a new scheme for conducting secure and anonymous online English auctions using a modified type of group signature. Trust is divided among three servers owned by separate companies to ensure anonymity and fairness. Our scheme solves the problems of the existing English auction schemes and has following characteristics: *unforgeability, anonymity, unlinkability, exculpability, coalition-resistance, verifiability, robustness, traceability, revocation, one-off registration, unskewability* and *unblockability*. Our scheme has comparable efficiency to the existing schemes for the enhanced security and privacy it provides.

1 INTRODUCTION

Online auctioning is now widely accepted as one of the premiere means to do business on the web. English auctions are the most common type of online auction employed by Internet auctioneers (e.g., eBay¹ and uBid²). Such auctions are used to sell various items from real estate to football tickets. An English auction allows one seller to offer an item for sale. Many potential buyers then submit bids for the item attempting to outbid each other. The winner is the bidder with the highest bid after a given time-out period where no bid higher than the current highest bid has been made. The winner must pay the seller an amount equal to the winning bid.

Since the participants are not physically present in an online auction, there exist many security concerns and opportunities for people to cheat. For example, a bidder might repudiate having made a bid or the seller doesn’t deliver the item. Furthermore, the Auctioneer could influence the auction in a manner inconsistent

with its rules (e.g., block bids). Security and privacy in electronic auctions has been covered in (Boyd and Mao, 2000; Franklin and Reiter, 1996; Naor *et al.*, 1991; Trevathan, 2005; Viswanathan *et al.*, 2000), and numerous “secure” auction schemes have been proposed. However, most of the schemes presented so far have been for *sealed bid* auctions (i.e., bids remain secret until the close of bidding). An English auction on the other hand is an *open bid* auction (i.e., everyone knows the values of the bids). This combined with the nature of the auctioning process makes English auctions more complicated than regular cryptographic auction schemes.

The timing of events in English auctions is much more critical than sealed bid auctions. As a result, this presents some unique security risks. An English auction requires a real-time link between the bidders and the Auctioneer. Frequent price quotes are issued to update bidders regarding the current highest bid. As bidders base their decisions on this information, its timeliness directly influences the auction. A corrupt Auctioneer could disadvantage certain bidders by delaying this information or by speeding up (skewing) the clock in order to close the auction early. Fur-

¹<http://www.ebay.com>

²<http://www.ubid.com>

thermore, the speed and ease of the bid submission process is significant, especially when an auction is nearing its end. A malicious Auctioneer could selectively block bids based on bidder identity and/or bid value.

(Stubblebine and Syverson, 1999) presented an English auction scheme that prevents the Auctioneer from closing the auction early and from blocking bids. However it does not protect a bidder's anonymity. Alternately, a scheme by (Omote and Miyaji, 2001) provides anonymity, but does not stop an Auctioneer from skewing its clock or blocking bids. We believe the short-comings of the existing schemes can be solved by basing the auction protocol on a modified group signature scheme.

The concept of group signatures was introduced by (Chaum and van Heyst, 1991). A group signature scheme allows members of a group to sign messages on behalf of the group, such that the resulting signature does not reveal the identity of the signer. Signatures can be verified with respect to a single group public key. Only a designated group manager is able to open signatures, and thus reveal the signer's identity. Due to these unique security characteristics, group signature schemes have recently been used as the basis for auction protocols (see (Trevathan *et al.*, 2005; Trevathan *et al.*, 2006)).

This paper presents a scheme for conducting online English auctions in a secure and anonymous manner. The new scheme solves the problems of the existing proposals while maintaining all of their features. The role of the Auctioneer is divided among two auction servers (owned by separate companies) to ensure that the correct timing of events is maintained and to prevent bid blocking. (see (Naor *et al.*, 1991).) Our scheme uses a group signature that is altered so that the role of the group manager is also divided among two independent auction servers. This allows for bid verification and protects a bidder's identity unless the two servers collude. In the case of a dispute (e.g., a bidder repudiates a bid), a court order can be used to reveal the bidder's identity and he/she can be permanently revoked from the auction proceedings. The scheme is flexible and allows the group signature to be updated as better techniques for group signatures become available. Our scheme offers comparable efficiency trade-offs for its enhanced security and privacy characteristics.

This paper is organised as follows: the remainder of this section discusses security issues inherent in English auctions and our contribution. Existing English auction schemes and their shortcomings are discussed in Section 2. The components of our new scheme are introduced in Section 3 and the auction protocol is described in Section 4. An informal security analysis of the new scheme is given in Section 5. Section 6 presents an efficiency comparison of the new scheme

and Section 7 provides some concluding remarks.

1.1 Fundamentals of Online English Auctions

There are four main activities in an online English auction:

Initialisation – The Auctioneer sets up the auction and advertises it i.e., type of good being auctioned, starting time, etc.

Registration – In order to participate in the auction, bidders must first register with the Auctioneer.

Bidding – A registered bidder computes his/her bid and submits it to the Auctioneer. The Auctioneer checks the bid received to ensure that it conforms with the auction rules.

Winner Determination – The Auctioneer determines the winner according to the auction rules. Online English auctions can terminate according to the following rules (see (Kumar and Feldman, 1998; Stubblebine and Syverson, 1999)):

1. *Expiration Time* - The auction closes at a predetermined expiration time.
2. *Timeout* - The auction closes when no bids higher than the current highest bid are made within a predetermined timeout interval.
3. *Combination of Expiration and Timeout* - The auction closes when there is a timeout after the expiration time.

1.2 Security Issues in Online English Auctions

The core security requirements for an English auction include:

Unforgeability - Bids must be unforgeable, otherwise a bidder can be impersonated.

Verifiability - There must be publicly available information by which all parties can be verified as having correctly followed the auction protocol. This should include evidence of registration, bidding and proof of the winner of the auction.

Exculpability - Neither the Auctioneer nor a legitimate bidder can forge a valid signature of a bidder.

Coalition-resistance - No coalition of bidders can frame an innocent bidder by fabricating a bid.

Robustness - The auction process must not be affected by invalid bids or by participants not following the correct auction protocol.

Anonymity - The bidder-bid relationship must be concealed so that no bidder can be associated or identified with the bid they submit.

One-time registration - Registration is a one-off procedure, which means that once a bidder has registered, they can participate in future auctions held by the Auctioneer.

Unlinkability - Bids are unlinkable within an auction, and also between plural auctions.

Traceability - Once a bidder has submitted a bid, they must not be able to repudiate having made it. Otherwise if a bidder wins and does not want to pay, they might deny that they submitted the winning bid. In this event the identity of the bidder who submitted the bid in question can be revealed.

Revocation - Malicious bidders can be easily revoked from all future auctions.

English auctions are open bid and the timely nature of the auction process therefore raises several further concerns. Due to the flexibility of closing rules for English auctions this introduces the following unique requirements:

Unskewability - The Auctioneer must not be able to alter the auction timing. For example, speed up its clock in an attempt to close the auction early, or slow the auction down to keep the bidding process active beyond the official timeout.

Unblockability - The Auctioneer cannot selectively block bids based on bid amount or the identity of the bidder.

Conditional bid cancellation - In online auctions using an expiration time, it is common for the auction to continue for days or weeks. In this situation a bidder might be reluctant to make such an open ended bid. Therefore depending on the closing rule and the stage of the auction it is desirable to allow bidders to conditionally cancel bids. Note that bidders should not be able to cancel bids when an auction is in a timeout stage and cancellation must only be done in strict accordance with the Auctioneer's bid cancellation policy.

2 EXISTING ENGLISH AUCTION SCHEMES

Discussions regarding security for English auctions can be found in (Kumar and Feldman, 1998; Trevathan *et al.*, 2005). Several "secure" English auction schemes have been proposed by (Lee *et al.*, 2001; Nguyen and Traore, 2000; Omote and Miyaji, 2001; Stubblebine and Syverson, 1999). The first scheme is due to (Stubblebine and Syverson, 1999). This scheme requires bidders to register with the Auctioneer. The Auctioneer must periodically timestamp the

auction proceedings with a *Notary* to prove to bidders that it is not skewing its clock. Bidders submit bids using a reverse hash chain and secret bid commitments. This is done to ensure that the Auctioneer cannot block bids, and that bidders are not able to repudiate bids. The auction proceedings are recorded on a public bulletin board that is readable by everyone, but can only be written to by the Auctioneer.

We have identified the following problems with this scheme:

1. There is no anonymity for the bidders.
2. Bids are linkable, meaning that the Auctioneer can create profiles about individual bidders and their bidding strategies.
3. All parties must trust the Notary. (i.e., to ensure the correct timing is maintained.)

(Omote and Miyaji, 2001) refine a scheme by (Nguyen and Traore, 2000) that uses a form of modified group signature (Ateniase *et al.*, 2000; Camenisch and Stadler, 1997; Chaum and van Heyst, 1991). This scheme allows a bidder to register once and participate in any number of auctions held by the Auctioneer. Bids are claimed to be unlinkable between different auctions, but linkable within a particular auction. This is achieved by requiring the bidder to calculate a new signature generation key prior to each auction.

In this scheme there are two managers responsible for conducting the auction. The *Registration Manager* (RM) secretly knows the correspondence of the bidder's identity and registration key. RM works as an identity escrow agency. The *Auction Manager* (AM) hosts the auction and prepares bidder's auction keys in each round.

We have identified the following problems with this scheme:

1. All bidders must update their keys between each round of auctioning, which is essentially equivalent to re-registering. Therefore, this negates the author's claims that registration is a one-off procedure.
2. AM can skew its clock and/or selectively block bids.
3. Revoking a bidder is inefficient as it requires AM to reissue new keys to all of the existing bidders.
4. (Lee *et al.*, 2001) describe a flaw in this scheme during the winner announcement stage. Here AM is able to erroneously inform any bidder that they have won without being publicly verifiable. Lee *et al.* propose a solution. However, this introduces several more bulletin boards and requires computations that are an order of magnitude slower.
5. Bids are linkable within a current auction, but unlinkable between plural auctions. The motivation

for this is stated as the auction participants gain utility in terms of entertainment from viewing the auction. For example, when there is a rally between two particular bidders, observers enjoy knowing how many bids a bidder has submitted.

With regard to the last point, it is our opinion, that in an anonymous auction scheme all bids (whether in the same auction or not) must be totally unlinkable. Observers can still see a rally, however, there is no need to know exactly whom the bids are coming from. Our scheme described in the next section, does not allow bids to be linked within the same auction or between plural auctions.

3 COMPONENTS OF OUR SCHEME

The auction has four parties:

A *Bidder*, who is interested in buying an item from a seller in an English auction.

An *Auction Manager* (AM), who organises the auction proceedings, accepts bids and determines the winner according to whoever has submitted the highest bid. To participate in an auction, a bidder presents his/her real identity to AM. AM issues the bidder with a token that allows him/her to register.

A *Registration Manager* (RM), who takes part in the protocol in order to complete the registration of a bidder, once a token has been obtained from AM. At the end of the protocol, the bidder obtains a secret key that enables him/her to generate signed bids in a proper format.

An *Auction Helper* (AH), who aids AM in accepting bids and determining the winner. AH is owned by a separate company and is tasked with ensuring that AM does not alter its clock or block bids.

The scheme uses a two-server trust approach that can be broken down into two subsystems: the *anonymity subsystem* and the *auction subsystem* (see Figure 1). The anonymity subsystem protects the anonymity of the bidders provided the AM and RM do not collude. The auction subsystem ensures the correct outcome of the auction as long as AM and AH do not collude. There is no trust assumed between RM and AH.

Each bidder, AM and AH are connected to a common broadcast medium with the property that messages sent to the channel instantly reach every party connected to it. The broadcast channel is public so

that everybody can listen to all information communicated via the channel, but cannot modify it. It is also assumed that there are private channels between RM and any potential bidders (who wish to join the auction proceedings).

3.1 Group Signatures

To join an auction, a bidder must first register with RM (who plays the role of a group manager in a group signature scheme). Once registered, a bidder can participate in the auction by signing bids using the group signature. Bids are submitted to an independent AM who runs the auction (with the help of AH which is explained later). AM (and AH) post the auction results on a publicly verifiable bulletin board.

One of the most efficient and popular proposals for group signature schemes is due to (Ateniese *et al.*, 2000). This is the group signature scheme that is used for the basis of our auction protocol. The (Ateniese *et al.*, 2000) group signature scheme informally works as follows:

Let $n = pq$ be an RSA modulus, where p and q are two safe primes (i.e., $p = 2p' + 1$, $q = 2q' + 1$, and p' , q' are also prime numbers). Denote by $QR(n)$, the set of quadratic residues - a cyclic group generated by an element of order $p'q'$. The group public key is $\mathcal{Y} = (n, a, a_0, y = g^x, g, h)$, where a, a_0, g, h are randomly selected elements from $QR(n)$. The secret key of the group manager is x .

To join the group, a user (bidder i) must engage in a protocol with the group manager (i.e., RM and AM) and receive a group certificate $[B_i, e_i]$ where $B_i = (a^{x_i}, a_0)^{1/e_i} \bmod n$ with e_i and x_i chosen from two integral ranges as defined in (Ateniese *et al.*, 2000). (x_i is only known to the user/bidder).

In order to sign a message/bid, m , the user/bidder has to prove possession of his member certificate $[B_i, e_i]$ without revealing the certificate itself. More precisely, the user/bidder computes:

$$T_1 = B_i y^w \bmod n, T_2 = g^w \bmod n,$$

$$T_3 = g^{e_i} h^w \bmod n SK(m)$$

where the value $SK(m)$, computed over a message m , indicates a signature of knowledge of the secret key x_i and the e_i th root of the first part of the representation of T_3 (in the implementation of our scheme, the exact signature generation and verification procedures will be presented).

In the case of a dispute, the group manager can open a signature that reveals the identity of the signer. This is due to the fact that the pair (T_1, T_2) is an El-Gamal encryption of the user's certificate (using the public key of the group manager). That is, the group manager can compute B_i , using $B_i = T_1 / (T_2)^x$.




Figure 1: The Auction Model

In certain circumstances users must be revoked from the group. For example, a membership expires or a user misbehaves. Reissuing keys to all existing group members is unwieldy and inefficient for a large group. Using a certificate revocation list to blacklist malicious bidders requires the verifier of the signature to check a list that is linear in the number of revoked users.

(Camenisch and Lysyanskaya, 2002) propose a scheme based on a dynamic accumulator that requires a member to prove that they have not been revoked. Informally, an *accumulator* is a method to combine a set of values into one short accumulator such that there is a short witness that a given value was incorporated into the accumulator. It is infeasible to find a witness for a value that is not in the accumulator. A *dynamic accumulator* allows values to be added and deleted from the accumulator at unit cost. By incorporating dynamic accumulators into a group signature scheme, revocation can easily be performed by deleting a member's value from the accumulator.

A user must check the accumulator prior to signing. This requires an online link between the group manager and the users. In terms of an auction, a bidder must check the accumulator each time they submit a bid. This is reasonable for English auctions, as there is a real-time communication link between the Auctioneer and bidders anyway.

The (Camenisch and Lysyanskaya, 2002) dynamic accumulator scheme can be defined as follows: A dynamic accumulator for a family of inputs $\{\mathcal{X}_1\}$ is a family of families of functions $\{\mathcal{F}_1\}$ with the following properties:

Efficient generation: *There is an efficient probabilistic algorithm \mathcal{G} that on input 1^k produces a random element f of \mathcal{F}_k . Moreover, along with f , \mathcal{G} also outputs some auxiliary information about f , denoted*

aux_f .

Efficient evaluation: *$f \in \mathcal{F}_k$ is a polynomial-size circuit that, on input $(u, k) \in \mathcal{U}_f \times \mathcal{X}_k$, outputs a value $v \in \mathcal{U}_f$, where \mathcal{U}_f is an efficiently-samplable input domain for the function f ; and \mathcal{X}_k is the intended input domain whose elements are to be accumulated.*

Quasi-commutative: *For all k , for all $f \in \mathcal{F}_k$ for all $u \in \mathcal{U}_f$ for all $x_1, x_2 \in \mathcal{X}_k$, $f(f(u, x_1), x_2) = f(f(u, x_2), x_1)$. If $\mathcal{X} = \{x_1, \dots, x_m\} \subset \mathcal{X}_k$, then by $f(u, \mathcal{X})$ we denote $f(f(\dots(f(u, x_1), \dots), x_m))$.*

Witness: *Let $v \in \mathcal{U}_f$ and $x \in \mathcal{X}_k$. A value $w \in \mathcal{U}_f$ is called a witness for x in v under f if $v = f(w, x)$.*

Addition: *Let $f \in \mathcal{F}_1$, and $v = f(u, \mathcal{X})$ be the accumulator so far. There is an efficient algorithm A to accumulate a given value $x' \in \mathcal{X}_1$. The algorithm outputs:*

1. $\mathcal{X}' = \mathcal{X} \cup \{x'\}$ and $v' = f(v, x') = f(u, \mathcal{X}')$;
2. w' which is the witness for $x \in \mathcal{X}$ in v' .

Deletion: *Let $f \in \mathcal{F}_1$, and $v = f(u, \mathcal{X})$ be the accumulator so far. There exist efficient algorithms \mathcal{D}, \mathcal{W} to delete an accumulated value $x' \in \mathcal{X}$. The functionality of the algorithms includes:*

1. $\mathcal{D}(aux_f, v, x') = v'$ such that $v' = f(u, \mathcal{X} \setminus \{x'\})$, and
2. $\mathcal{W}(w, x, x', v, v') = v'$ such that $f(w', x) = v'$, where $x \in \mathcal{X}$ and $f(w, x) = v$.

The (Camenisch and Lysyanskaya, 2002) dynamic accumulator scheme is based on the strong RSA assumption and accumulates prime numbers (i.e., the primes used for the membership certificates in (Ateniese *et al.*, 2000) group signature scheme). The scheme also provides a proof that a committed value was accumulated (we will omit these details). The construction of a dynamic accumulator where the domain of accumulated values consists of prime numbers, is as follows:

- F_k is the family of functions that correspond to exponentiating modulo-safe prime products drawn from the integers of length k . Choosing $f \in F_k$ amounts to choosing a random modulus $n = pq$ of length k , where $p = 2p' + 1$, $q = 2q' + 1$, and p, p', q, q' are all prime. We will denote f corresponding to modulus n and domain $\mathcal{X}_{A,B}$ by $f_{n,A,B}$.

- $\mathcal{X}_{A,B}$ is the set $\{e \in \text{primes} : e \neq p', q' \wedge A \leq e \leq B\}$, where A and B can be chosen with arbitrary polynomial dependence on the security parameter k , as long as $2 < A$ and $B < A^2$. $\mathcal{X}'_{A,B}$ is (any subset of) the set of integers from $[2, A^2 - 1]$ such that $\mathcal{X}_{A,B} \subseteq \mathcal{X}'_{A,B}$.

- For $f = f_n$, the auxiliary information aux_f is the factorisation of n .

- For $f = f_n$, $\mathcal{U}_f = \{u \in QR_n : u \neq 1\}$ and $\mathcal{U}'_f = Z_n^*$.

- For $f = f_n$, $f(u, x) = u^x \bmod n$. Note that $f(f(u, x_1), x_2) = f(u(x_1, x_2)) = u^{x_1 x_2} \bmod n$.

- Update of the accumulator value. Adding a value \tilde{x} to the accumulator value v can be done as $v' = f(v, \tilde{x}) = v^{\tilde{x}} \bmod n$. Deleting a value \tilde{x} from the accumulator is as follows: $\mathcal{D}((p, q), v, \tilde{x}) = v^{\tilde{x}-1 \bmod (p-1)(q-1)} \bmod n$.

- Update of a witness. Updating a witness u after \tilde{x} has been added can be done by $u' = f(u, \tilde{x}) = u^{\tilde{x}}$. In case, $\tilde{x} \neq x \in \mathcal{X}_k$ has been deleted from the accumulator, the witness u can be updated as follows. By the extended GCD algorithm, one can compute the integers a, b such that $ax + b\tilde{x} = 1 \bmod n$ and then $u' = \mathcal{W}(u, x, \tilde{x}, v, v') = u^{bv^a}$.

4 THE AUCTION PROTOCOL

This section describes the auction protocol. A high level view of the protocol is given in Figure 2. Lines depict communication between parties while the dashed circles indicate stages in the protocol. Lines that pass through the dashed circles are communications that are performed during the particular stage.

4.1 Setup

Most activities of this stage need to be performed only once (in order to establish the auction proceedings). Let $\lambda_1, \lambda_2, \gamma_1$, and γ_2 be some lengths, Λ, Γ be some integral ranges, and $\mathcal{H}(\cdot)$ be a collision-resistant hash function. RM sets up the group public key and his secret key by performing the following steps:

1. Chooses two safe primes p and q (i.e., $p = 2p' + 1$ and $q = 2q' + 1$, where p' and q' are prime numbers) and sets the RSA modulus $n = pq$
2. Chooses random elements $a, a_0, g, h \in QR(n)$

3. Chooses a secret element $x \in_R Z_{p'q'}^*$ and sets $y = g^x \bmod n$
4. Publishes the group public key as $\mathcal{Y} = (n, a, a_0, y, g, h)$
5. Creates the public modulus n for the accumulator, chooses a random $u \in QR_n$ and publishes (n, u)
6. Set up (empty for now) public archives E_{add} for storing values that correspond to added users and E_{delete} for storing values that correspond to deleted users

4.2 Registration

A user submits a request to AM to participate in the auction proceedings. AM verifies the identity of the requestor, and issues a token that is verifiable by RM. The user then takes part in a protocol with RM, in order to obtain his/her secret key and a certificate of membership in the auction proceedings. Note that the token does not carry the real identity of the bidder. All communication between RM and the owner of a token is authenticated and recorded. The protocol between a new bidder i , and RM is as follows (checks in which values are chosen from proper intervals, the user knows discrete logarithms of values, etc. are omitted):

1. Bidder i selects random exponents x'_i, r and sends $C_1 = g^{x'_i} h^r \bmod n$ to the RM
2. RM checks that $C_1 \in QR(n)$. If this is the case, RM selects random values α_i, β_i and sends them to bidder i
3. Bidder i computes $x_i = 2^{\lambda_1} + (\alpha_i x'_i + \beta_i \bmod 2^{\lambda_2})$ and sends to RM the value $C_2 = a^{x_i} \bmod n$
4. RM checks that $C_2 \in QR(n)$. If this is the case, RM selects a random $e_i \in \Gamma$ and computes $B_i = (C_2 a_0)^{1/e_i} \bmod n$ then sends the membership certificate $[B_i, e_i]$ to bidder i (note that $B_i = (a^{x_i} a_0)^{1/e_i} \bmod n$)
5. Bidder i verifies that $a^{x_i} a_0 = B_i^{e_i} \bmod n$
6. Add the current u to the bidder's membership certificate. Update u : $u = f_n(u, e_i)$. Update E_{add} : store e_i there
7. Verify that $f_n(u_i, e_i) = u_i^{e_i} = u$

RM creates a new entry in the membership table and stores bidder i 's membership certificate $[B_i, e_i]$ and a transcript of the registration process in this location.

4.3 Setup - before each auction

AM organises the auction (i.e., advertising and calls for auction). AM posts information to the bulletin




Figure 2: The Auction Protocol

board regarding the auction including the auction id (which uniquely identifies the auction), the reserve price (minimum winning price that will be accepted), the auction starting time and the auction closing rules.

4.4 Bidding

Using a membership certificate $[B_i, e_i]$, a bidder can generate anonymous and unlinkable group signatures on a bid m . m contains the auction id and the amount of the bid (i.e., $m = id \parallel \text{bid value}$). Bidder i submits a bid m to both AM and AH signed using his/her secret key.

Update Membership - Prior to submitting a bid, a bidder must check if there have been any changes to the group (i.e., new bidders have been added, or other bidders have been revoked). If this is the case, a bidder must perform a membership update. This is done as follows:

An entry in the archive is called “new” if it was entered after the last time bidder i performed an update.

- Let y denote the old value of u
- For all new $e_j \in E_{add}$, $u_i = f(u_i, \prod e_j) = u_i^{\prod e_j}$ and $y = y^{\prod e_j}$
- For all new $e_j \in E_{delete}$, $u_i = W(u_i, e_i, \prod e_j, y, u)$ (Note that as a result $u = f(u_i, e_i)$)

Sign Bid - In order to generate a signature on a message/bid, m , bidder i performs the following:

- Chooses a random value w and computes:

$$T_1 = B_i y^w \text{ mod } n, \quad T_2 = g^w \text{ mod } n,$$

$$T_3 = g^{e_i} h^w \text{ mod } n$$

- Chooses r_1, r_2, r_3, r_4 (randomly) from predetermined intervals and computes:

$$(a) \quad d_1 = T_1^{r_1} / (a^{r_2} y^{r_3}), \quad d_2 = T_2^{r_1} / (g^{r_3}), \quad d_3 = g^{r_4}, \quad \text{and } d_4 = g^{r_1} h^{r_4} \text{ (all in mod } n),$$

$$(b) \quad c = \mathcal{H}(g \parallel h \parallel y \parallel a_0 \parallel a \parallel T_1 \parallel T_2 \parallel T_3 \parallel d_1 \parallel d_2 \parallel d_3 \parallel d_4 \parallel m),$$

$$(c) \quad s_1 = r_1 - c(e_i - 2^{\xi_1}), \quad s_2 = r_2 - c(x_i - 2^{\lambda_1}), \\ s_3 = r_3 - ce_i w, \quad \text{and } s_4 = r_4 - cw \text{ (all in } \mathbb{Z}).$$

- In addition to T_1 , T_2 , and T_3 the bidder computes the values $C_e = g^e h^{r_1}$, $C_u = u h^{r_2}$, and $C_r = g^{r_2} h^{r_3}$ and sends them to AM, with random choices $r_1, r_2, r_3 \in_R \mathbb{Z}_{[n/4]}$

- The output is

$$(c, s_1, s_2, s_3, s_4, r_1, r_2, r_3, r_4, T_1, T_2, T_3, C_e, C_u, C_r)$$

Prove Membership/Verify Bid - AM and AH check the validity of the bidder’s signature using the group’s public key \mathcal{Y} . A bid of the correct form is considered to be valid and is included in the auction (i.e., posted on the bulletin board). An invalid bid is discarded. There are two copies of the bid on the bulletin, one posted by AM and the other posted by AH. AM and AH verify the signature on the bid as follows:

- Compute (all in mod n):

$$c' = \mathcal{H}(g \parallel h \parallel y \parallel a_0 \parallel a \parallel T_1 \parallel T_2 \parallel T_3 \parallel (a_0^c T_1^{(s_1 - c 2^{\xi_1})}) / (a^{s_2 - c 2^{\lambda_1}} y^{s_3}) \parallel (T_2^{s_1 - c 2^{\xi_1}}) / (g^{s_3}) \parallel T_2^c g^{s_4} \parallel T_3^c g^{s_1 - c 2^{\xi_1}} h^{s_4} \parallel m)$$

2. AM, AH and the bidder engage in a protocol to prove membership (see (Camenisch and Lysyanskaya, 2002) for details)
3. Accept the signature if and only if $c = c'$, and the parameters s_1, s_2, s_3, s_4 lie in the proper intervals

Bid Cancellation - If a bidder desires to cancel a bid, they must send a copy of the bid they wish to cancel and a CANCEL message signed using his/her group key to both AM and AH. Upon receiving the CANCEL message, AM and AH check the bidder's signature on the message using the group's public key \mathcal{Y} . If the signature is valid, AM and AH then check what stage the auction is in. If the auction close rule is currently in an expiration time stage, AM and AH each post a message to the bulletin stating that the particular bid has been cancelled. If the auction is currently in a timeout stage, the CANCEL message is discarded and the bid remains in effect.

4.5 Winner Determination

Once the auction has closed, AM and AH then determine the auction outcome according to which bidder has made the highest bid. The winning bidder can produce a copy of the signed bid as evidence that they have won.

4.6 Traceability

In the event of a dispute, RM (with the help of AM) can open the signature on a bid to reveal which bidder is the original signer. This process is as follows:

1. Check the signature's validity via the verification procedure
2. Recover B_i (and thus the identity of bidder i) as $B_i = T_1/T_2^x \bmod n$

RM then checks the registration transcripts, and determines the token associated with this certificate. AM, who knows the relation between tokens and real identities, can determine the identity of the bidder. Note that in our scheme, revealing the identity of a bidder does not reveal any information about his/her past bids.

4.7 Revocation

When a bidder has been caught breaking the auction rules, they can be permanently revoked from the auction proceedings by cancelling the bidder's ability to sign future bids. To achieve this, the bidder's prime number used in his/her membership certificate is not included when the dynamic accumulator is updated. This can be done as follows: Retrieve e_i which is the

prime number corresponding to the bidder's membership certificate. Update u : $u = \mathcal{D}(\psi(n), u, e_i)$. Update E_{delete} : store e_i there.

5 SECURITY

This section provides an informal security analysis of the online English auction scheme presented in this paper based on the characteristics described in Section 1.2.

Unforgeability - Only bidders that are members of the group are able to sign messages on behalf of the group. This is due to the unforgeability of the underlying group signature.

Anonymity - Given a valid signature $(c, s_1, s_2, s_3, s_4, T_1, T_2, T_3)$ identifying the actual signer is computationally difficult. Determining which bidder with certificate $[B_i, e_i]$ has signed a bid, requires deciding whether the three discrete logarithms $\log_y T_1/B_i$, $\log_g T_2$, and $\log_g T_3/g^{e_i}$ are equal. This is assumed to be infeasible under the decisional Diffie-Hellman assumption, and thus anonymity is guaranteed. Note that in our auction, RM can figure out the certificate associated with each signature, but cannot determine the identity of the bidder associated with this certificate.

Unlinkability - Deciding if two signatures $(c, s_1, s_2, s_3, s_4, T_1, T_2, T_3)$ and $(\tilde{c}, \tilde{s}_1, \tilde{s}_2, \tilde{s}_3, \tilde{s}_4, \tilde{T}_1, \tilde{T}_2, \tilde{T}_3)$ were computed by the same bidder is computationally hard (with the same argument as for anonymity).

Exculpability - Neither a bidder nor AM, AH and/or RM can sign on behalf of another bidder. This is because the secret key x_i , associated to user i is computationally hidden from RM. RM, at most, can learn $a^{x_i} \bmod n$, which cannot help him to learn the exponent x_i (since the discrete logarithm over the safe composite modulo n , is difficult).

Coalition-resistance - This is due to the following theorem: (Ateniese *et al.*, 2000) Under the strong RSA assumption, a group certificate $[B_i = (a^{x_i} a_0)^{1/e_i} \bmod n, e_i]$ with $x_i \in \Lambda$ and $e_i \in \Gamma$ can be generated only by the group manager provided that the number K of certificates the group manager issues is polynomially bounded.

Verifiability - All bids (including signatures) are posted to the public bulletin, therefore all parties can verify the auction outcome.

Robustness - Invalid bids will not be posted to the bulletin board. Moreover, malicious bidders will be revoked from the system, and thus cannot affect the auction outcome.

Traceability - RM is always able to open a valid signature and, with the help of AM, identify the signer

Table 1: Comparison of English auction schemes

	SS99	OM001	Our Scheme	TX03
Registration	1 exp.	480 mul.	30 exp.	2 exp.
Signing	1 exp.	240 mul.	25 exp.	17 exp.
Verification	1 exp.	320 mul.	21 exp.	16 exp.
Revocation	N/A	$O(\ell)$	$O(1)$	$O(1)$

of the bid.

Revocation - Bidders can be easily revoked from the future auctions if they have broken the auction rules. See theorem 2 in (Camenisch and Lysyanskaya, 2002).

One-time registration - Once a bidder has received a signature generation key, they are free to participate in future auctions.

Unskewability - AH observes AM’s clock (and vice versa) therefore any clock skews will not go unnoticed. AM’s clock can be trusted as long as both AM and AH do not collude.

Unblockability - A bidder must submit his/her bids to both AM and AH, who post the bid on the bulletin board. If either tries to block a bid, then only one confirmation of the bid will be posted to the bulletin board which will indicate that one of the parties has blocked a bid. Bids cannot be blocked unless AM and AH collude.

Conditional bid cancellation - Bidders can conditionally cancel bids by sending a CANCEL message to AM and AH as long as the auction is not in a time-out stage.

6 EFFICIENCY

This section discusses the efficiency considerations of the new scheme. We contrast our approach with the existing English auction schemes. Table 1 shows the amount of work performed during each major stage of the auction in terms of the number of modular exponentiations (exp) or multiplications (mul) required. The schemes compared include: (Stubblebine and Syverson, 1999) (SS99), (Omote and Miyaji, 2001) (OM01), our scheme, and (Tsudik and Xu, 2003) (TX03). ((Tsudik and Xu, 2003) is an alternate implementation of our approach.)

The registration, signing and verification procedures for SS99 are relatively efficient. However, SS99 do not protect a bidder’s identity, nor do they discuss revocation issues. To incorporate revocation into this scheme, it is likely that the registration procedure would have to be repeated between auctions. Furthermore, SS99 do not address the issue of one-time registration. Once again bidders would have to repeat

the registration process for each auction they want to participate in.

OM01 is significantly less efficient than SS99. OM01 does not address bid cancellation whereas SS99 does. Furthermore, OM01 does not prevent the Auctioneer from skewing its clock. However, OM01 protects a bidders identity and addresses one-time registration. The cost of one-time registration in OM01 is issuing new keys to bidders between auctions, which is essentially equivalent to re-registering. The revocation method in OM01 is tied in with the one-time registration mechanism and therefore must also be repeated between each auction. To revoke a bidder requires the Auctioneer to perform work proportional to $O(\ell)$ where ℓ is the number of bidders.

In contrast, our scheme has the most practical one-time registration procedure. That is, once a bidder has registered, there is no work required to retain membership other than regularly checking the accumulator. We address bid cancellation, clock-skewing and privacy concerns. To revoke a bidder, the Auctioneer only has to update the accumulator. Bidders must check the accumulator value prior to each bid which is a constant operation. Our auction scheme can also be implemented using TX03 which has significant efficiency gains.

The efficiency of our scheme is comparable to the existing proposals. First of all our scheme has an enhanced set of security requirements that are much more comprehensive. Furthermore, our scheme clearly has the most efficient revocation method. In addition, we have the most practical one-time registration procedure.

7 CONCLUSIONS

This paper presented a scheme for conducting secure and anonymous online English auctions. Such a scheme is vital for protecting the security and anonymity of participants who engage in online auctioning. The timeliness of information and verifiability of the Auctioneer’s actions is critical in an online English auction. We have shown that the existing “secure” English auction schemes are inadequate for the task. The scheme by (Stubblebine and Syver-

son, 1999) does not provide anonymity for the bidders and requires all parties to trust a public Notary. The scheme by (Omote and Miyaji, 2001) does not prevent an Auctioneer from skewing his/her clock or from blocking bids.

In direct contrast, our scheme solves all of the problems of the existing schemes and has a more comprehensive set of security requirements. We use a group signature to provide verification of bids and to protect the identities of bidders. The group signature is modified so that the identity of a bidder is divided among two separate parties (i.e., the anonymity subsystem). The role of the Auctioneer is also divided among two parties to prevent clock-skewing and bid-blocking (i.e., the auction subsystem). The scheme has comparable efficiency to the existing proposal for its enhanced security and privacy characteristics. The efficiency and security of the scheme rests with the underlying group signature scheme used. Our approach offers the client flexibility in choosing from any group signature scheme. The scheme offers efficient one-time registration and revocation procedures that are clearly better suited to handling multiple auctions than existing proposals.

REFERENCES

- Ateniese, G., Camenisch, J., Joye, M. and Tsudik, G. (2000). A practical and provably secure coalition secure coalition-resistant group signature scheme in *Advances in Cryptology - Proceedings of CRYPTO 2000*, vol. 1880 of *Lecture Notes in Computer Science*, Springer-Verlag, 255-270.
- Ateniese, G., Song, D. and Tsudik, G. (2002). Quasi-Efficient Revocation of Group Signatures, in *Proceedings of Financial Cryptography*, vol. 2357 of *Lecture Notes in Computer Science*, Springer-Verlag, 183-197.
- Boyd, C. and Mao, W. (2000). Security Issues for Electronic Auctions, Technical Report, *Hewlett Packard, TR-HPL-2000-90*.
- Camenisch, J. and Lysyanskaya, A. (2002). Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials, in *Advances in Cryptology - Proceedings of CRYPTO 2002*, vol. 2442 of *Lecture Notes in Computer Science*, Springer-Verlag, 61-76.
- Camenisch, J. and Stadler, M. (1997). Efficient Group Signature Scheme for Large Groups, in *Advances in Cryptology - Proceedings of CRYPTO '97*, vol. 1294 of *Lecture Notes in Computer Science*, Springer-Verlag, 410-424.
- Chaum, D. and van Heyst, E. (1991). Group Signatures, in *Advances in Cryptology - Proceedings of EURO-CRYPT'91*, vol. 547 of *Lecture Notes in Computer Science*, Springer-Verlag, 257-265.
- Franklin, M. and Reiter, M. (1996). The Design and Implementation of a Secure Auction Service, *IEEE Transactions on Software Engineering*, vol. 22, 302-312.
- Kumar, M. and Feldman, S. (1998). Internet Auctions, in *Proceedings of the Third USENIX Workshop on Electronic Commerce*, 49-60.
- Lee, B., Kim, K. and Ma, J. (2001). Efficient Public Auction with One-time Registration and Public Verifiability, in *International Conference on Cryptology in India - Proceedings of INDOCRYPT 2001*, vol. 2247 of *Lecture Notes in Computer Science*, Springer-Verlag, 162-174.
- Naor, M., Pinkas, B. and Sumner, R. (1999). Privacy Preserving Auctions and Mechanism Design, in *The 1st ACM Conference on Electronic Commerce*, 129-139.
- Nguyen, K. and Traore, J. (2000). An On-line Public Auction Protocol Protecting Bidder Privacy, in *Proceedings of ACSIP 2000 - Australasian Conference on Information Security and Privacy*, vol. 1841 of *Lecture Notes in Computer Science*, Springer-Verlag, 427-442.
- Omote, K. and Miyaji, A. (2001). A Practical English Auction with One-Time Registration, in *Proceedings of ACSIP 2001 - Australasian Conference on Information Security and Privacy*, vol. 2119 of *Lecture Notes in Computer Science*, Springer-Verlag, 221-234.
- Stubblebine, S. and Syverson, P. (1999). Fair On-line Auctions Without Special Trusted Parties, in *Proceedings of Financial Cryptography 1999*, vol. 1648 of *Lecture Notes in Computer Science*, Springer-Verlag, 230-240.
- Tsudik, G. and Xu, S. (2003). Accumulating Composites and Improved Group Signing, in *Advances in Cryptology - Proceedings of ASIACRYPT 2003*, vol. 2894 of *Lecture Notes in Computer Science*, Springer-Verlag, 269-286.
- Trevathan, J. (2005). Security, Anonymity and Trust in Electronic Auctions, *Association for Computing Machinery Crossroads*, Spring Edition, 3-9, vol. 11.3.
- Trevathan, J., Ghodosi, H. and Read, W. (2005). Design Issues for Electronic Auctions, in *2nd International Conference on E-Business and Telecommunication Networks*, 340-347.
- Trevathan, J., Ghodosi, H. and Read, W. (2006). An Anonymous and Secure Continuous Double Auction Scheme, in *39th International Hawaii Conference on System Sciences*, 125(1-12).
- Viswanathan, K., Boyd, C. and Dawson, E. (2000). A Three Phased Schema for Sealed Bid Auction System Design, *Proceedings of ACSIP 2000 - Australasian Conference on Information Security and Privacy*, vol. 1841 of *Lecture Notes in Computer Science*, Springer-Verlag, 412-426.

Chapter 5

Fraudulent Auctioning Practices

Fraudulent auctioning practices are not easy to protect against using the cryptographic methods of previous chapters. This section presents three papers. The first paper outlines the problems of undesirable behaviour in auctions (see Section 5.1). The remaining papers present methods to detect shill bidding in online English auctions. The first of these papers concentrates on the situation where there is only one shill (see Section 5.2). The next paper extends the detection mechanisms to incriminate multiple shills working in collusion with each other (see Section 5.3).

5.1 Undesirable and Fraudulent Behaviour in Online Auctions

Auction participants can behave in an undesirable and fraudulent manner in an attempt to gain an advantage at the expense of rivals. For example, a bidder might seek to suppress the price by bid sniping, or the seller could introduce fake bids to inflate the price. In addition, an outsider or rival seller can lure away bidders by directly offering them better deals, or a malicious seller can auction misrepresented or non-existent items. This conduct is a problem as it results in market failure, thereby inhibiting the usefulness of online auctions as an exchange medium. While cryptography has been used to provide security in terms of bid authentication and privacy, there is no documented means to prevent many of the aforementioned problems. This paper investigates undesirable and fraudulent behaviour in online auctions. The following practices are examined: bid shielding, shill bidding, bid sniping, siphoning and selling non-existent or misrepresented items. The characteristics of such behaviour are described, and how to identify it in an auction. The paper also provides recommendations for recourse against undesirable and fraudulent participants.

“Undesirable and Fraudulent Behaviour in Online Auctions” [9] is published as a short paper in the International Conference on Security and Cryptography 2006 (SECRYPT). SECRYPT 2006 was held in Setúbal, Portugal and was organised by INSTICC. Sponsors included IBM, Setúbal Polytechnic Institute, IEEE Systems, Man and Cybernetics (SMC) Society, and the International Association for Cryptographic Research (IACR). The major goal of this conference was to bring together researchers and developers from academia and industry working in areas related to e-business to stimulate a debate with a special focus on telecommunications networks. The conference received 326 submissions in total, with contributions from 53 different countries. To evaluate each submission, a double blind paper evaluation method was used: each paper was reviewed by at least two internationally known experts from the Program Committee, and more than 95% of the papers had 3 or more reviews. In the end, 98 papers were selected

to be published and presented as full papers. This corresponds to a 30% acceptance ratio. An additional 68 papers were published and presented as short papers, cumulatively amounting to a 51% acceptance ratio. This paper is available online via citeseer ¹ and is referenced by DBLP ².

Publication in SECRYPT 2006 was made possible by funding from the School of Mathematical and Physical Sciences.

¹<http://citeseer.ist.psu.edu/>

²<http://www.informatik.uni-trier.de/~ley/db/>

UNDESIRABLE AND FRAUDULENT BEHAVIOUR IN ONLINE AUCTIONS

Jarrold Trevathan

*School of Mathematical and Physical Sciences
James Cook University
Email: jarrod.trevathan@jcu.edu.au*

Wayne Read

*School of Mathematical and Physical Sciences
James Cook University
Email: wayne.read@jcu.edu.au*

Keywords: Bid shielding, shilling, sniping, siphoning, non-existent/misrepresented items.

Abstract: Online auctions are a popular means for exchanging items over the Internet. However, are many inherent security and fairness concerns. Participants can behave in an undesirable and fraudulent manner in an attempt to gain an advantage at the expense of rivals. For example, a bidder might seek to suppress the price by bid sniping, or the seller could introduce fake bids to inflate the price. In addition, an outsider or rival seller can lure away bidders by directly offering them better deals, or a malicious seller can auction mis-represented or non-existent items. This conduct is a problem as it results in market failure, thereby inhibiting the usefulness of online auctions as an exchange medium. While cryptography has been used to provide security in terms of bid authentication and privacy, there is no documented means to prevent many of the aforementioned problems. This paper investigates undesirable and fraudulent behaviour in online auctions. We examine the following practices: bid shielding, shill bidding, bid sniping, siphoning and selling non-existent or misrepresented items. We describe the characteristics of such behaviour and how to identify it in an auction. We also provide recommendations for recourse against undesirable and fraudulent participants.

1 INTRODUCTION

Online auctions are one of the most popular destinations on the web. Buyers and sellers can exchange items amongst a worldwide audience from the comfort and privacy of their own homes. Participants remain largely anonymous and can bid in any manner they desire. However, this freedom comes at the expense of new security risks and fairness concerns.

By behaving in an undesirable or fraudulent manner, one party is able to gain at the expense of another. For example, bidders can use practices such as *bid shielding* and *bid sniping* to keep the price low. Alternately, *shill bidding* is a strategy which a seller may pursue, to artificially inflate the auction price. In addition, *siphoning* is a tactic employed by an outsider, who is seeking to profit from an auction by offering bidders a cheaper, identical item. Finally, a seller might attempt to auction a *non-existent or misrepresented item*.

Auction security has been previously discussed in (Franklin and Reiter, 1996; Stubblebine and Syverson, 1999; Trevathan *et al*, 2005). Cryptographic methods have been proposed to solve many of these security issues (see e.g., (Franklin and Reiter, 1996;

Viswanathan *et al*, 2000; Trevathan *et al*, 2006)). However, cryptographic solutions are generally limited to bid authentication and privacy. Protecting auction participants from the aforementioned problems is a much harder task (i.e., a bidder's bidding strategy or misrepresented goods). Furthermore, none of the proposed models in literature specifically address the auction format used in online auctions.

This paper discusses undesirable and fraudulent practices in online auctions. We describe the characteristics of such behaviour, and show how auction participants can identify if they are a victim. We also provide recommendations/options for what to do when such behaviour is encountered. We show how behaviours are related, and when an auction exhibits one form of undesirable behaviour, the other forms soon manifest to counter balance it. This inevitably leads to market failure, and retards the usefulness of online auctions as an exchange medium.

This paper is organised as follows: Section 2 describes the basic format and rules of a typical online auction. Sections 3 through 7 discuss each major form of undesirable auctioning behaviour. Each of these sections includes a description of individual behavioural characteristics, remedies/recourse for victims



Figure 1: An Example Online Auction




Figure 2: Bid Shielding

and relationships between certain types of behaviour. Section 8 provides some concluding remarks.

2 ONLINE AUCTIONS

There are many types of auction (e.g., Vickrey, CDA, etc.). The most popular type of auction is the English auction. In an English auction, bidders outbid each other in an attempt to win an item. The winner is the bidder with the highest bid. English auctions are commonly employed in online auctions such as those offered by eBay¹ and ubid²

Online auctions differ to traditional auctions in that the auction ends after a given period of time. Figure 1 illustrates a typical online auction (such as that offered by eBay). Time flows from left to right. Bidders can only submit bids between the auction start and end times. Each bid must be for an amount that is greater than the current highest bid. When the auction terminates, the winner is the bidder with the highest bid (in this case \$75).

eBay and ubid offer a mechanism for automatically bidding on a bidder's behalf. This is referred to as the *proxy bidding system* on eBay, and the *bid butler* on ubid. A bidder is only required to enter the maximum price they are willing to pay. The bidding software will then automatically outbid any other bid until the maximum, potentially saving the bidder money. Such mechanisms remove the need for a bidder to be constantly watching an auction for bidding activity.

¹<http://www.ebay.com>

²<http://www.ubid.com>

3 BID SHIELDING

Bid shielding typically involves several bidders working in collusion with each other, or a bidder with access to multiple accounts. The first bidder enters a bid for the amount they are willing to pay for an item. A second bidder then immediately enters an excessively high bid in an attempt to deter other bidders from continuing to bid. As the auction is drawing to a close, the second bid is retracted. When this occurs, the next highest bid remaining (i.e., the first bid) then becomes the winning bid.

Figure 2 illustrates a bid shielding scenario. Initially regular bidders make a few bids (\$10 and \$15 in this case). The first colluding bidder enters his/her bid for \$20. This is referred to as the *shielded bid*. The second colluding bidder then enters a bid for \$250, which is well above the current bidding range, and is unlikely to be outbid. This is known as the *shielding bid*.

In the example, the shielding bid deters regular bidders from bidding again. Prior to the end of the auction, the second colluding bidder retracts the shielding bid for \$250. The winning bid then falls back to the shielded bid for \$20. As the bid is retracted near the end of the auction, other bidders do not have time to respond. By this stage, most regular bidders have lost interest in the auction anyway.

Bid shielding generally cannot be accomplished as shown in the example due to the presence of software bidding agents (i.e., eBay's proxy bidding system). As soon as the shielding bid is entered, the bidding history won't show this as a bid for \$250. Instead, this will be listed as the minimal amount required to out-



Figure 3: Bid Shielding in Auctions with Automated Bidding Agents

bid the shielded bid (e.g., \$21, as \$1 is the minimum increment). Rival bidders will then be inclined to continue bidding, which in turn drives up the shielding bid's value.

Bid shielding in the presence of automated bidding agents can be achieved by using two shielding bids. Figure 3 illustrates this scenario. As in the previous example, a shielded bid for \$20 is submitted. A shielding bid for \$250 is entered followed immediately by another shielding bid for \$260. As the second shielding bid outbids the first, the bidding history immediately reflects the value of the first shielding bid (i.e., \$250), thus giving the appearance of a high price. Near the end, **both** of the shielding bids are retracted. The winning bid then drops to \$20 (i.e., the shielded bid).

This sort of behaviour disadvantages the seller in terms of suppressing the price. Honest bidders are also disadvantaged as they are forced out of the auction.

There are no mechanisms in place to prevent bid shielding. Most auctioneers keep a record of the number of bid retractions made by a bidder. However, the purpose for the record keeping is actually to deter bidders from renegeing on winning bids (i.e., win and then later decide they don't want the good). Retractions that occur before an auction terminates are also typically deemed less suspicious than those that occur afterwards. Furthermore, keeping records on retraction rates is largely useless as bidders can simply register under different aliases.

Another type of price suppressing behaviour is referred to as *pooling* or *bid rigging*. Two or more bidders collude, and agree not to bid up the price. However, this attack is only effective if the number of bidders is small, or the majority of bidders are in collusion.

4 SHILL BIDDING

Shill bidding (or shilling) is the act of introducing fake bids into an auction on the seller's behalf in order to artificially inflate the price of an item. Bidders who engage in shilling are referred to as *shills*. To

win the item, a legitimate bidder must outbid a shill's price. If one of the shills accidentally wins, then the item is re-sold in a subsequent auction. Shill bidding is a problem as it forces legitimate bidders to pay significantly more for the item.

In March 2001, a U.S. federal grand jury charged three men for their participation in a ring of fraudulent bidding in hundreds of art auctions on eBay (see (Schwartz and Dobrzynski, 2002)). The men created more than 40 user IDs on eBay using false registration information. These aliases were used to place fraudulent bids to artificially inflate the prices of hundreds of paintings they auctioned on eBay.

The men hosted more than 1,100 auctions on eBay from late 1998 until May 2000, and placed *shill bids* on more than half of those auctions. The total value of the winning bids in all auctions which contained shill bids exceeded approximately \$450,000. The total value of the shill bids in these auctions exceeded approximately \$300,000 (equivalent to 66%).

To be effective, a shill must comply to a particular strategy which attempts to maximise the pay-off for the seller. This section provides an insight into the general behaviour of shills. It describes a shill's characteristics and strategies, presents examples of shill behaviour in an auction, and discusses shill detection techniques.

4.1 Shill Mindset

The main goal for shilling is to artificially inflate the price for the seller beyond the limit that legitimate bidders would otherwise pay to win the item. The pay-off for the seller is the difference between the final price and the uninflated price. A shill's goal is to lose each auction. A shill has an infinite budget. If the shill wins, the item will have to be re-auctioned. Resale of each item costs the seller both money and effort thereby eroding the possible gains from shilling.

The shill faces a dilemma for each bid they submit. Increasing a bid could marginally increase the revenue for the seller. However, raising the price might also result in failure if it is not outbid before the auction terminates. The shill must decide whether to take the deal or attempt to increase the pay-off.




Figure 4: Aggressive Shill Bidding

On the contrary, a bidder's goal is to win. A bidder has a finite budget and is after the lowest price possible. Increasing a bid for a legitimate bidder decreases the money saved, but increases the likelihood of winning.

4.2 Shill Characteristics and Strategies

A shill has the following characteristics:

1. A shill usually bids exclusively in auctions only held by one particular seller, however, this alone is not sufficient to incriminate a bidder. It may be the case that the seller is the only supplier of an item the bidder is after, or that the bidder really trusts the seller (usually based on the reputation of previous dealings).
2. A shill tends to have a high bid frequency. An aggressive shill will continually outbid legitimate bids to inflate the final price. Bids are typically placed until the seller's expected payoff for shilling has been reached, or until the shill risks winning the auction (e.g., near the termination time or during slow bidding).
3. A shill has few or no winnings for the auctions participated in.
4. It is advantageous for a shill to bid within a small time period after a legitimate bid. Generally a shill wants to give legitimate bidders as much time as possible to submit a new bid before the closing time of the auction.
5. A shill usually bids the minimum amount required to outbid a legitimate bidder. If the shill bids an amount that is much higher than the current highest bid, it is unlikely that a legitimate bidder will submit any more bids and the shill will win the auction.
6. A shill's goal is to try and stimulate bidding. As a result, a shill will tend to bid more closer to the beginning of an auction. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the end of an auction is risky as the shill could accidentally win.

The most extreme shill bidding strategy is referred to as *aggressive shilling*. An aggressive shill continually outbids everyone thereby driving up the price as much as possible. This strategy often results in the shill entering many bids.

In contrast, a shill might only introduce an initial bid into an auction where there has been no prior bids with the intent to stimulate bidding. This kind of behaviour is a common practice in both traditional and online auctions. However, most people typically do not consider it fraudulent. Nevertheless it is still shilling, as it is an attempt to influence the price by introducing spurious bids.

This is referred to as *benign shilling* in the sense that the shill does not continue to further inflate the price throughout the remainder of the auction. A benign shill will typically make a "one-off" bid at or near the very beginning of the auction.

Regardless of the strategy employed, a shill will still be a bidder that often trades with a specific seller but has not won any auctions.

Another factor that affects a shill's strategy is the value of the current bid in relation to the reserve price. For example, once bidding has reached the reserve price, it becomes more risky to continue shilling. However, this is conditional on whether the reserve is a realistic valuation of the item that all bidders share.

4.3 Shill Bidding Examples

Figure 4 illustrates an example auction with aggressive shilling. The shill aggressively outbids a legitimate bid by the minimal amount required to stay ahead, and within a small time period of the last bid. The shill bids force the other bidders to enter higher bids in order to win. The shill does not win the auction despite the high number of bids;

Figure 5 illustrates an example auction with benign shilling. Initially no bids have been made. A shill bid is entered for \$10 to try stimulate bidding. After seeing that there is some demand for the item, other bidders eventually submit bids for the item. The shill does not enter any further bids.




Figure 5: Benign Shill Bidding

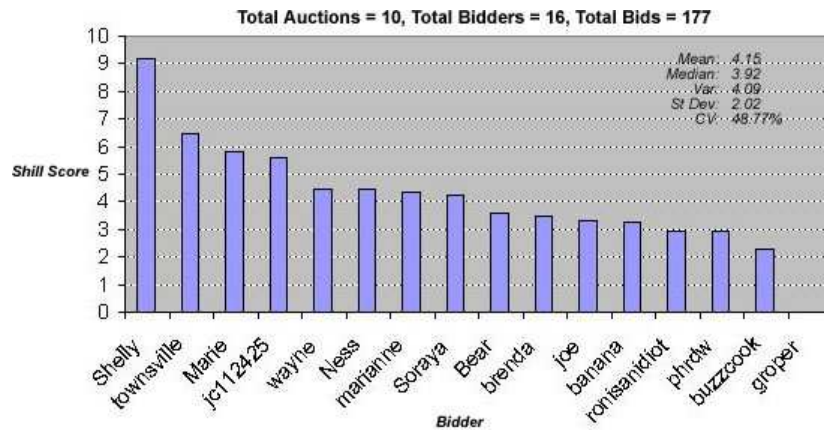


Figure 6: An example of the Shill Score when run on auction simulations containing one aggressive shill

4.4 Shill Detection

There is often much confusion regarding what constitutes shill behaviour. Bidding behaviour that might seem suspicious could in fact turn out to be innocent. Furthermore, a shill can engage in countless strategies. This makes it difficult to detect shill bidding. While the online auctioneers monitor their auctions for shilling, there is no academic material available on proven shill detection techniques.

(Wang *et al*, 2002) suggest that listing fees could be used to deter shilling. Their proposal charges a seller an increasing fee based on how far the winning bid is from the reserve price. The idea is to coerce the seller into stating their true reserve price, thereby eliminating the economic benefits of shilling. However, this method is untested and does not apply to auctions without reserve prices.

We are developing techniques to detect shill bidding (Trevathan and Read, 2005). Our method examines a bidder's bidding behaviour over several auctions and gives them a shill score to indicate the degree of suspicious behaviour they exhibit. Bidders who engage in suspicious price inflating behaviour will rate highly, whereas those with more regular bidding behaviour will rate low. A bidder can examine other bidder's shill scores to determine whether they

wish to participate in an auction held by a particular seller.

The shill score has been tested using simulated auctions involving real world people. To facilitate testing we implemented an online auction server (see (Trevathan and Read, 2006)). Several types of tests have been conducted. The first type involves auctioning fake items to real bidders. Each bidder is allocated a random amount of money, which they use to bid in the auction. Even though bidders don't actually receive the item, these tests manage to recreate the mental drive and desire to win. Winners are excited and often boastful after a hard fought auction. One person (namely the author) is tasked with being a shill in order to stimulate bidding. The shill's goal is to increase the price as much as possible, without actually winning the auction.

When the shill score is used on these auctions, it clearly identifies the shill bidder. It also exonerates innocent bidders that bid in a regular manner. The shill scores for a series of tests is given in Figure 6. In this case, the bidder known as Shelly is the shill bidder. Shelly engaged in aggressive shilling behaviour and consequently has a shill score that is over nine. This is clearly much higher than the other bidders. The results from these tests are thus far encouraging.

We have also conducted similar auctions using real



Figure 7: Bid Sniping

items (e.g., bottles of wine and collector's edition playing cards), where the winner was required to pay real money. In these settings, bidders are more cautious. However, the shill was still able to influence the auction proceedings and also was detected by the shill score. (Note that all shilling victims were fully reimbursed!) Furthermore, the shill score has been tested using commercial auction data and simulations involving automated bidding agents.

5 BID SNIPING

Bid sniping is another undesirable type of bidding behaviour. A bidder who employs a sniping strategy is referred to as a *sniper*. A sniper will only bid in the closing seconds of an auction, thereby denying other bidders time to react. This essentially prevents the sniper from being outbid.

Figure 7 illustrates the mechanics of bid sniping. Regular bidders enter their bids as normal. In the closing seconds, the sniper enters a bid for the minimum required to win (i.e., \$41). The bid entered by a sniper is referred to as a *sniper bid*. None of the other bidders have time to outbid the sniper bid, and therefore the sniper wins.

Sniping behaviour is the exact opposite of shilling. A sniper's goal is to win the auction for the lowest price. Whereas a shill's goal is not to win and to inflate the price. Sniping is often used as a preventative measure against shilling. A sniper may not be able to prevent shilling occurring during an auction. However, the sniper can prevent themselves from being shilled.

Sniping disadvantages regular bidders in that they are denied the opportunity to respond to the sniper bid. Bidders are typically frustrated when they realise that sniping has occurred. This is especially the case when a bidder has observed an auction for a long period of time, only to be beaten in the closing seconds. The seller is also potentially disadvantaged by sniping, as the sniper bid does not stimulate rival bidding, that might have occurred, had other bidders been able to respond.

Sniping is permitted on eBay, although its use is discouraged. Instead eBay recommends that a bidder should only place a single bid at his/her maximum valuation using the proxy bidding system. Despite this recommendation, sniping is rampant, and is now considered as a natural part of the online auctioning experience. (Shah *et al*, 2002) performed a study into the amount of sniping in 12,000 eBay auctions. Their results showed for the majority of auctions, a significant fraction of bidding occurs in the closing seconds.

A *sniping agent* is a software bidding agent that follows a sniping strategy. The sniping agent constantly monitors an auction, and waits until the last moment to bid. Many companies now exist such as Bidnapper.com³, ezsniper⁴ and Auction Sniper⁵ which offer sniping agents for use on eBay auctions.

uBid auctions differ to eBay in that auctions terminate using a timeout session. Once the ending time of an auction has been reached, the auction is extended for ten minutes for each bid received. This limits the effectiveness of sniping, but can lead to a show down between snipers. As a result the auction can run for a lot longer than the seller anticipated. The seller can enforce a maximum extension limit to prevent the auction continuing indefinitely. However, this results in the auction being essentially the same as a normal auction without a timeout session.

The only preventative measure for a bidder against sniping, is to "out-snipe" the sniper. However, this often results in there being *multiple snipers* in an auction. This behaviour leads to failure of the English auctioning process. If everyone engaged in sniping, the auction would essentially become a *sealed bid* auction. In a sealed bid auction, bidders submit their bids secretly during a bidding round. At the close of bidding, the Auctioneer determines the winner. English auctions on the other hand are *open bid*, and allow bidders to bid multiple times. In a traditional offline English auction, sniping cannot occur. Sniping is a feature unique to online auctions. Sniping behaviour blurs the boundaries of an online auction between the

³<http://www.bidnapper.com>

⁴<http://www.ezsniper.com>

⁵<http://www.auctionsniper.com>

type of auction it is and the rules that govern it.

6 SIPHONING

Siphoning (or bid siphoning) refers to the situation where an outsider observes an auction and contacts bidders offering them an identical item at a better price. The outsider is referred to as a *siphoner*, and is said to “siphon” bids from the auction. The siphoner benefits in that he/she does not incur any of the costs involved with organising and advertising an auction.

Siphoning disadvantages the Auctioneer through lost revenue. That is, the siphoner does not have to pay the Auctioneer to list/advertise an item. Siphoning also disadvantages the seller whose auctions are being siphoned. This is in a form of price undercutting (i.e., the siphoner offers the item at a better price), and reduces the demand for the seller’s items. A bidder who does business with a siphoner loses the protection offered by the auction, and exposes themselves to fraud (e.g., misrepresented or non-existent items).

Siphoning may also be used in conjunction with shilling. When a shill bid accidentally wins, the seller of the item can contact the next highest bidder, and directly offer them the item. This saves the seller the time and expense of re-auctioning the item.

Consider the following scenario: A seller is auctioning off a traditional Japanese sword. A legitimate bidder enters a bid for \$1000. The seller then enters a shill bid for \$1200. The legitimate bidder refrains from bidding and the shill bid wins. Later on the legitimate bidder is approached by the seller. The seller claims that the winner is a *dead beat*⁶, or has backed out of the deal due to financial or other personal reasons. The seller then offers the item to the bidder at, or near, the bidder’s price of \$1000.

Siphoning combined with shilling disadvantages both the bidders and the Auctioneer. Bidders are forced into paying an inflated price due to the shill bids. The siphoning component this time does not affect the seller, but rather the Auctioneer. This is because the seller does not have to repeat the auction and is denying the Auctioneer revenue from listing fees. The seller in effect has “siphoned” bids from his/her own future auctions.

Siphoning is impossible to detect by the Auctioneer alone. Instead bidders must report the behaviour once it has happened to them. However, most bidders are aware of what siphoning is, and probably wouldn’t recognise that they have been siphoned. Furthermore, if a bidder receives a better price, or an item they re-

⁶A bidder that has won an auction and fails to make payment.

ally want, then they would not see such behaviour as undesirable. In this case they are unlikely to report it.

There is no clear law regarding siphoning. The only advice to bidders is to decline communication with anyone that contacts them outside of the official channels. In addition, if you were not the winner of an auction, don’t accept an item from the seller after the auction has ended.

7 NON-EXISTENT OR MISREPRESENTED ITEMS

A dubious seller might attempt to auction a non-existent item, or misrepresent an item. In the first instance, the seller accepts payment from the buyer but doesn’t deliver the item. In the second, the seller misrepresents the item by advertising it as something it isn’t, or delivers an item of lesser value. In the shill case described in Section 4, the men misrepresented paintings as being significant works when they were inexpensive replicas (see (Schwartz and Dobrzynski, 2002)).

To ensure that an item conforms to its description, eBay recommends that buyers study the seller’s photos. However, this is unsatisfactory as the seller can simply copy pictures from a legitimate item, and post these for the non-existent/misrepresented item.

eBay also recommends asking the seller questions regarding the item. However, a seller may simply lie. Furthermore, some bidders might be reluctant to ask questions as they desire anonymity. For example, if the bidder has a high profile, they might want to conceal the fact that they are bidding. This might occur in the case where the bidder feels that their privacy would be compromised in some manner (e.g., reveal that they have a fetish for an item), change other bidders’ perceptions of the item (e.g., stimulate unwanted bidding), or influence the seller to raise the reserve price.

Another recommendation is to check seller feedback. eBay has a system where buyers and sellers can leave feedback regarding their dealings with each other. A party to a transaction can rate their experience as either good, neutral or bad. An individual is given a rating based on the feedback received. However, feedback ratings are not a reliable measure of an individual’s integrity, and feedback can be falsely generated using multiple bidder accounts.

Online auctioneers offer a dispute resolution procedure where a buyer can initiate action against a seller if an item was not received, or if the item is different from what was expected. If a seller constantly does not deliver items, the Auctioneer can revoke his/her account. However, the seller can simply re-register under a different alias.

Insurance can be offered to buyers for items not delivered. eBay has created a successful off-shoot business, PayPal⁷, for guaranteeing online transactions. However, launching an insurance claim can be an arduous process, and might not fully compensate the victim. Furthermore, it can require the compensated funds to be spent by participating in future auctions. In addition, it is only a solution that can be used **after** fraud has occurred.

Escrow fraud is an emerging threat to online auctions. A fraudster sets up a fake escrow service for paying a seller for an item. When the winning bidder wires the money to the escrow agency, the agency vanishes along with the payment. Neither the seller receives the payment, nor does the winner receive the item. The following are typical characteristics of an escrow scam:

1. Check for poor grammar on the escrow site.
2. Although site may look authentic, it is usually copied from a legitimate site such as Escrow.com⁸ and Auctionchex⁹.
3. There are obvious give-aways in the **Terms** page, which is generally stolen from another site.
4. A site will often leave hints of what its previous incarnation was - especially if they've just changed domain names recently.
5. Be wary if the seller insists on using a specific escrow site. Sellers don't usually press for escrow, buyers do.

Bidders must be wary of the escrow service they use, and not do business with unknown or un-trusted escrow vendors.

8 CONCLUSIONS

Online auctions are susceptible to undesirable and fraudulent behaviour. Such behaviour is designed to influence the auction in a manner that either favours a bidder, the seller, or an outsider seeking to profit from the auction. This results in market failure, and reduces the usefulness of online auctions as an exchange medium.

This paper investigates undesirable and fraudulent trading behaviour, and discusses its implications for online auctioning. We show how to identify such behaviour and give recommendations for recourse against undesirable and fraudulent participants.

Bid shielding is a practice employed by one or more bidders to suppress bidding and keep the price

low. There is no means to protect against bid shielding, other than to exclude users with a high bid retraction rate. However, a bidder can use a fake name to re-register.

Alternately shill bidding is a practice employed by the seller to artificially inflate the price by introducing spurious bids. Shill bidding is strictly forbidden by commercial online auctions, and is a prosecutable offence. There is limited material on shill bidding and prevention/detection techniques. We are presently developing a method to detect shill bidders by giving them a ranking called a shill score. A bidder can use the shill score to decide whether they want to participate in an auction held by a particular seller.

Bid sniping is a strategy employed by a bidder to prevent being outbid. A sniper submits a bid during the closing seconds of an auction, thereby denying other bidders time to react. Many bidders engage in sniping behaviour in an attempt to prevent themselves from being shilled. Sniping is permitted in online auctions, but is discouraged. Commercial sniping agents are available that engage in sniping behaviour on the bidder's behalf. Sniping can be reduced by enforcing a timeout limit which extends the auction by several minutes for each new bid received after the auction's termination time. The prominence of sniping behaviour raises concerns regarding the effectiveness of online auctions to conduct English auctions according to traditional rules.

Siphoning refers to the situation where an outsider observes an auction and offers an identical item to the bidders at a lower price. The siphoner avoids all of the costs associated with conducting an auction, and effectively profits from the seller. Siphoning is often used in conjunction with shilling, and scams involving non-existent or misrepresented items. Siphoning behaviour should be reported to the Auctioneer, however, there is little more that can be done.

A malicious seller can offer non-existent or misrepresented items to profit from unsuspecting buyers. It is recommended that bidders inspect photos and ask questions. However, a seller can post fake photos, lie in response to questions, and a buyer might want to remain anonymous. eBay's feed back rating is a mechanism for gauging the integrity of a buyer/seller, however, it is dubious at best. Insurance can reimburse victims of fraud, but is not a complete solution. Escrow fraud is emerging as a threat to insurance-based fraud prevention measures and is difficult to identify. The best advice is for a buyer to be wary of the goods they are bidding for.

The popularity of online auctions and the overwhelming number of flawless transactions, is a testament to the success of online auctioning. Nevertheless, some participants are disadvantaged by undesirable or fraudulent behaviour. Preventing such behaviour is a difficult task and existing solutions are

⁷<http://www.paypal.com>

⁸<http://www.escrow.com>

⁹<https://auctionchex.com>

largely inadequate.

Buying items from online auctions is the same as buying items anywhere online. A buyer cannot physically inspect the merchandise as in a “bricks and mortar” store. The buyer is forced to rely on the item’s description. Distances between buyers and sellers can be vast. This makes it hard to police transactions that go awry, especially when buyers and sellers cross political and cultural boundaries.

The best recommendation is *caveat emptor*. Don’t deal with unknown sellers, or sellers who reside in countries which do not have strict enforcement of international commerce laws. Ensure that you thoroughly research the item you are bidding for.

REFERENCES

- Franklin, M. and Reiter, M. (1996). The Design and Implementation of a Secure Auction Service, *IEEE Transactions on Software Engineering*, vol. 22, 302-312.
- Schwartz, J. and Dobrzynski, J. (2002). 3 men are charged with fraud in 1,100 art auctions on eBay, in *The New York Times*.
- Shah, H., Joshi, N. and Wurman, P. (2002). Mining for Bidding Strategies on eBay, in *SIGKDD’2002 Workshop on Web Mining for Usage Patterns and User Profiles*.
- Stubblebine, S. and Syverson, P. (1999). Fair On-line Auctions Without Special Trusted Parties, in *Proceedings of Financial Cryptography 1999*, vol. 1648 of *Lecture Notes in Computer Science*, Springer-Verlag, 230-240.
- Trevathan, J., Ghodosi, H. and Read, W. (2005). Design Issues for Electronic Auctions, in *Proceedings of the 2nd International Conference on E-Business and Telecommunication Networks (ICETE)*, 340-347.
- Trevathan, J., Ghodosi, H. and Read, W. (2006). An Anonymous and Secure Continuous Double Auction Scheme, in *Proceedings of the 39th International Hawaii Conference on System Sciences (HICSS)*, 125(1-12).
- Trevathan, J. and Read, W. (2005). Detecting Shill Bidding in Online English Auctions, *Technical Report*, James Cook University.
- Trevathan, J. and Read, W. (2006). RAS: a system for supporting research in online auctions, *ACM Crossroads*, ed. 12.4, 23-30.
- Viswanathan, K., Boyd, C. and Dawson, E. (2000). A Three Phased Schema for Sealed Bid Auction System Design, *Proceedings of ACSIP 2000 - Australasian Conference on Information Security and Privacy*, vol. 1841 of *Lecture Notes in Computer Science*, Springer-Verlag, 412-426.
- Wang, W., Hidvegi, Z. and Whinston, A. (2002). Shill Bidding in Multi-round Online Auctions, in *Proceedings of the 35th Hawaii International Conference on System Sciences (HICSS)*.

5.2 Detecting Shill Bidding in Online English Auctions

This paper presents an algorithm to detect the presence of shill bidding in online auctions. It observes bidding patterns over a series of auctions, providing each bidder a score indicating the likelihood of their potential involvement shill behaviour. The algorithm has been tested on data obtained from a series of realistic simulated auctions as well as real data acquired from commercial online auctions. Results show that the algorithm is able to prune the search space required to detect which bidders are likely to be shills. This has significant practical and legal implications for commercial online auctions (such as eBay) where shilling is considered a major threat. While shilling is recognised as a problem, presently there is little or no established means of defence against shills. This paper presents a framework for a feasible solution, which acts as a detection mechanism and a deterrent. The shill score is extended in Section 5.3.

“Detecting Shill Bidding in Online English Auctions” [13] has been invited for publication as a book chapter in “Social and Human Elements of Information Security: Emerging trends and counter measures”, published by the Idea Group. This paper is available online via citeseer³.

This paper has resulted in media articles by JCU Outlook and ABC Learning. Interest has also been shown by The New Scientist and the Australian Scientist magazines. Dr Ick Jai Lee arranged commercial auction data and provided some initial comments. Dr Lee specialises in Data Mining.

³<http://citeseer.ist.psu.edu/>

Detecting Shill Bidding in Online English Auctions

Jarrod Trevathan

School of Maths, Physics and IT

James Cook University

Email: jarrod.trevathan@jcu.edu.au

Wayne Read

School of Maths, Physics and IT

James Cook University

Email: wayne.read@jcu.edu.au

Abstract—The popularity of online auctioning has grown meteorically since the late 90s. This attraction is due to their convenience, low cost and ability to reach large (even worldwide) audiences. However, despite the advantages there are many problems inherent in auctioning online. One such problem is shill bidding where spurious bids are introduced into an auction to drive up the final price for the seller, thereby defrauding legitimate bidders. This paper presents an algorithm to detect the presence of shill bidding in online auctions. It observes bidding patterns over a series of auctions, providing each bidder a score indicating the likelihood of their potential involvement shill behaviour. The algorithm has been tested on data obtained from a series of realistic simulated auctions as well as real data acquired from commercial online auctions. Our results show that the algorithm is able to prune the search space required to detect which bidders are likely to be shills. This has significant practical and legal implications for commercial online auctions (such as eBay) where shilling is considered a major threat. While shilling is recognised as a problem, presently there is little or no established means of defence against shills. This paper presents a framework for a feasible solution, which acts as a detection mechanism and a deterrent.

I. INTRODUCTION

Shill bidding is the act of introducing fake bids into an auction on behalf of the seller to artificially inflate the price of an item. The seller can either: 1) register as a bidder under a false identity or 2) be in collusion with one or more of the bidders. Bidders who engage in shilling are referred to as ‘shills’. To win the item on auction a legitimate bidder must outbid a shill’s price. If one of the shills accidentally wins, then the item is re-sold in a subsequent auction. Shill bidding is a problem as it forces legitimate bidders to bid against false bids hence paying significantly more for the item.

The most common type of auction is the English auction. In an English auction a seller offers an item for sale and many bidders attempt to outbid each other to obtain the item. The winner is the bidder with the highest bid after a given time-out period. English auctions are particularly susceptible to shill bidding practices.

The advent of online auctions such as eBay¹ and ubid² have made shill bidding much more exploitable. This is because it is relatively simple for a seller to register under many aliases and operate in rings with impunity. Furthermore, as bidders are not physically present it becomes much easier for a shill to anonymously influence the bidding process.

¹<http://www.ebay.com>

²<http://www.ubid.com>

This paper examines shill behaviour in the online setting and presents an algorithm to detect the presence of shill bidders in English auctions. The algorithm examines bidding information across several auctions and produces a score indicating the likelihood that a bidder is engaging in shill behaviour.

The algorithm is applied to real auction data where shills are present. We believe this is the first formal (documented) attempt to detect shills using a conformatory analysis of bidding behaviour. This paper shows that the algorithm can effectively reduce the search space required to detect shills. This enables the implementation of exploratory analysis using data mining techniques.

This paper is organised as follows: Section II contains a brief exploration of related work that has been conducted into shilling. Online English auctions are examined in Section III and Section IV discusses general shill behaviour in such auctions. Strategies for detecting shill bidders are addressed in Section V. Section VI presents an algorithm to detect shill bidding. The algorithm’s performance is described in Section VII using real auction data collected from both simulated and commercial auctions.

II. RELATED WORK

This section explores some related work that has been conducted into shilling. At present there is limited coverage on how to detect shill bidding.

eBay has been involved in many legal disputes where bidders/sellers have been accused of shilling (see [Schwartz and Dobrzynski 2002]). eBay has clear rules regarding shill bidding behaviour in their auctions³. Their policy clearly outlines undesirable bidder behaviour and the penalties for shill bidding. The regular process for a bidder who suspects that they have been shilled is to contact eBay, who then investigates the incident. eBay does not state exactly what factors they use to determine whether shilling has occurred nor how to detect which bidders are shills.

[Shah et al. 2002] use data mining techniques to produce evidence of shilling. Their work used data from approximately 12,000 commercial auctions looking for associations between bidders and sellers. Bidders (or groups of bidders) that participated frequently in auctions held by particular sellers were deemed suspect. However, the authors’ state that their analysis is very limited in that it only looks for simple associations.

³<http://pages.ebay.com/help/policies/seller-shill-bidding.html>

They suggest that a much more thorough analysis must be performed using complex associations which consider a wider range of shill behaviour.

There are also companies who offer data mining techniques to detect fraud in online auctions⁴. However, like eBay, these companies have not made their techniques public.

[Wang et al. 2002], [Wang et al. 2004] discuss an approach which attempts to deter shilling in the first place. The Auctioneer is allowed to use fees to make shilling unprofitable for the seller. In auctions using a reserve price, a seller is charged an increasing fee based on how far the winning price is from the reserve price.

Aspects regarding the economic theory of shilling have been examined by [Barbaro and Bracht 2005], [Kauffman and Wood 2003].

III. ONLINE ENGLISH AUCTIONS

This section describes the operation of online English auctions, which is the auction model used in this paper. Formally an English auction can be defined as an ascending-price, open-bid auction. Each bid submitted must be higher than the current highest bid. The minimal amount required to outbid others is usually a percentage of the current highest bid. The value of the current highest bid is available to all parties, along with the auction timing. The winner is the bidder with the highest bid when the auction terminates.

There are four main stages in an online English auction:

Initialisation: The Auctioneer sets up the auction and advertises it (i.e., description of goods, starting time, etc).

Registration: In order to participate in an online auction, a bidder must first register with the Auctioneer. During this stage a bidder establishes his/her identity (usually via credit card) and obtains a unique bidder id which they use to bid with.

Bidding: A registered bidder computes his/her bid and submits it to the Auctioneer. The Auctioneer checks the bid received to ensure conformity with the auction rules. Bids are posted on a bulletin board, where all bidders are able to obtain a real-time price quote and view the auction proceedings.

Winner Determination: At the close of bidding the Auctioneer determines the winner. Online English auctions can terminate according to the following rules (see [Kumar and Feldman 1998], [Stubblebine and Syverson 1999]):

- 1) *Expiration Time* - The auction closes at a predetermined expiration time.
- 2) *Timeout* - The auction closes when no bids higher than the current highest bid are made within a predetermined timeout interval.

3) *Combination of Expiration and Timeout* - The auction closes when there is a timeout after the expiration time.

The structure of bid information available to participants is as follows:

$$\langle b_{id}, price, time \rangle$$

where b_{id} is a unique bidder id, $price$ is the value the bidder is willing to pay, and $time$ is the time when the bidder submitted a bid. It is assumed that a bidder is only after a single quantity of an item.

In online auctions it is common to allow bidders to cancel (or retract) bids. Although eBay does not formally allow bid retraction, they permit it in exceptional cases such as typographical errors (e.g., the bidder entered \$99.9 instead of \$9.99), or if the item has radically changed in its description. Furthermore, auctions may last days or weeks and some bidders might be reluctant to make such opened-ended bids. For these reasons, bid retraction mechanisms are desirable in online auctions even though their use is discouraged.

In this paper it is assumed that a bidder cannot cancel a bid. When a bid has been submitted it is valid until it is beaten by a higher valued bid. If a bid is submitted that is less than the current highest bid, it is discarded.

IV. SHILL BEHAVIOUR

This section provides an insight into the general behaviour of shills. It describes a shill's mindset, characteristics and strategies, and presents an example of shill behaviour in an auction.

A. Shill Mindset

The main goal for shilling is to artificially inflate the price for the seller beyond what legitimate bidders would otherwise require to win the item. The pay-off for the seller is the difference between the final price and the uninflated price.

A shill's goal is to lose each auction. A shill has an infinite budget. If the shill wins, the item will have to be re-auctioned. Resale of each item costs the seller both money and effort thereby eroding the possible gains from shilling.

The shill faces a dilemma for each bid they submit. Increasing a bid could marginally increase the revenue for the seller. However, raising the price might also result in failure if it is not outbid before the auction terminates. The shill must decide whether to take the deal or attempt to increase the payoff.

On the contrary, a bidder's goal is to win. A bidder has a finite budget and is after the lowest price possible. Increasing a bid for a legitimate bidder decreases the money saved, but increases the likelihood of winning.

B. Shill Characteristics and Strategies

The following outlines typical shill characteristics assuming we are using an auction that terminates at a set expiration time:

- 1) A shill tends to bid exclusively in auctions only held by one particular seller, however, this alone is not sufficient to incriminate a bidder. It may be the case that the seller is the only supplier of an item the bidder is after, or

⁴<http://www.virtualgold.com>

that the bidder really trusts the seller (based on previous dealings).

- 2) A shill tends to have a high bid frequency. An aggressive shill will continually outbid legitimate bids to inflate the final price. A shill typically will bid until the seller’s expected payoff for shilling has been reached. Or until the shill risks winning the auction (e.g., near the termination time or during slow bidding).
- 3) A shill has few or no winnings for the auctions participated in.
- 4) It is advantageous for a shill to bid within a small time period after a legitimate bid. Generally a shill wants to give legitimate bidders as much time as possible to submit a new bid before the closing time of the auction.
- 5) A shill usually bids the minimum amount required to outbid a legitimate bidder. If the shill bids an amount that is much higher than the current highest bid, it is unlikely that a legitimate bidder will submit any more bids and the shill will win the auction.

It is common behaviour for legitimate bidders trying to conserve their money to also only bid the minimal amount. In situations where a bidder’s valuation is higher than the current bid, a bidder is more likely to outbid by more than the minimum amount.

- 6) A shill’s goal is to try and stimulate bidding. As a result, a shill will tend to bid more near the beginning of an auction. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the end of an auction is risky as the shill could accidentally win.

The most extreme shill bidding strategy is referred to here as *aggressive shilling*. An aggressive shill continually outbids everyone thereby driving up the price as much as possible. This strategy often results in the shill entering many bids.

In contrast, a shill might only introduce an initial bid into an auction where there has been no prior bids with the intent to stimulate bidding. This kind of behaviour is a common practice in both traditional and online auctions. However, most people typically do not consider it fraudulent. Nevertheless it is still shilling, as it is an attempt to influence the price by introducing spurious bids.

This is referred to as *benign shilling* in the sense that the shill does not continue to further inflate the price throughout the remainder of the auction. A benign shill will typically make a “one-off” bid at or near the very beginning of the auction.

Regardless of the strategy employed, a shill will still be a bidder that often trades with a specific seller but has not won any auctions.

Another factor that affects a shill’s strategy is the value of the current bid in relation to the reserve price. For example, once bidding has reached the reserve price it becomes more risky to continue shilling. This is conditional on whether the reserve is a realistic valuation of the item that all bidders share.

TABLE I
AN EXAMPLE AUCTION WITH ONE SHILL

Bid #	b_{id}	Price	Time
15	b_1	\$ 33	20:03
14	b_2	\$ 32	12:44
13	b_1	\$ 31	12:42
12	b_2	\$ 26	5:05
11	b_1	\$ 25	5:02
10	b_2	\$ 21	2:47
9	b_3	\$ 20	2:45
8	b_2	\$ 15	1:07
7	b_1	\$ 14	1:05
6	b_2	\$ 9	0:47
5	b_3	\$ 8	0:45
4	b_2	\$ 6	0:20
3	b_3	\$ 5	0:19
2	b_2	\$ 2	0:06
1	b_1	\$ 1	0:05

C. Shill Example

Table I illustrates an example auction with three bidders. Each bidder is denoted as b_1 , b_2 and b_3 respectively. Bidders b_1 and b_3 are legitimate whereas b_2 is a shill. b_2 engages in aggressive shill behaviour by outbidding a legitimate bid by the minimal amount required to stay ahead and within a small time period of the last bid. b_2 ’s bids force the other bidders to enter higher bids in order to win. If b_2 was not participating in this auction b_1 would have only needed to pay \$21 in order to win. Instead b_2 caused b_1 to pay \$33, thus the shill has inflated the price by \$12.

In this example, b_2 exhibits the typical shill behaviour described above. This is evidenced by: 1. High frequency of bids (i.e., b_2 has submitted more bids than both the other bidders); 2. Has not won the auction despite the high number of bids; 3. Quick to bid after a legitimate bidder; and 4. Only bids the minimal amount to stay in front.

V. SHILL DETECTION

This section describes how to detect shill bidding in online English auctions. It discusses the approach used in this paper and contrasts other less reliable proposals for shill detection.

As stated in Section I a shill can either be a bidder who is in collusion with the seller, or the seller can directly introduce fake bids of his/her own by registering as a bidder. For all intents and purposes, the distinction does not make any difference for detecting the presence of shill behaviour.

The solution used in this paper observes bidding patterns over a series of auctions for a particular seller looking for the shilling behaviour outlined in Section IV-B. The goal is to obtain statistics regarding the bidding patterns of bidders and deduce a measure called a *shill score*, which indicates the likelihood of a bidder/seller engaging in shill behaviour.

It has been suggested by [Wang et al. 2002] that this approach is inadequate as shills can change their strategies to adapt to any detection algorithm. While this is true, our algorithm targets core strategies that a shill follows. A shill

who deviates too far from these characteristics is less effective and won't significantly alter the auction outcome for the seller. Examining bidding behaviour acts as both a detection mechanism and a deterrent to shill bidders. To avoid detection, a shill must behave like a normal user, which in effect stops them from shilling.

There are also other characteristics possessed by shills. However, the use of these as a means of detection is dubious. For example, when a shill accidentally wins, the seller usually re-sells the item in a subsequent auction. The most obvious solution for detecting shilling is to observe if the same item is resold in a subsequent auction by the seller. This requires assigning a unique *item id* to each item sold by the seller. If the same *item id* is detected twice (i.e., the item has been resold by the seller), this provides some indication that shilling is taking place.

The problem with this approach is that it is only relevant for situations where a shill accidentally wins the item. That is, shilling is not detected in the case where a legitimate bidder has won and paid an inflated price. Furthermore, the seller might attempt to re-list the item under a new *item id*. This then allows them to re-sell the item without raising suspicion. For these reasons, tracking resold items is not a sound method for detecting shilling.

Another commonly used detection method exploits the property that shills tend to be within close geographical proximity to each other. There are two main reasons for this. Firstly, collusion among bidders typically occurs amongst friends who live/operate near each other. This is because the costs of communication and coordinating the shilling process are less than over long distances. Secondly, if the seller has several aliases, these aliases will be registered in the same location.

However, geographical proximity alone is an unreliable indicator of shill behaviour. First of all, the decreasing costs of communication hardware has made coordinating shill bidding over long distances much more affordable. Secondly, geographical information cannot be deduced from the raw auction data (as can the other characteristics) and involves cross-referencing users with the registration database, which raises anonymity concerns.

One proposed shill detection method examines the source IP address of a bidder. This is to prevent a shill from bidding using the same computer by logging in under different aliases. However, this is unreliable if more than one legitimate bidder must use the computer to bid (e.g., a computer in an Internet Cafe or a public library). Furthermore, this raises definite anonymity and privacy issues with regard to the bidders.

Other detection techniques examine seller feedback ratings and other statistics provided by eBay's feedback system. The feedback system allows buyers and sellers to report on their dealings with each other, which becomes an indicator of an individual's honesty and reliability. In many cases shill bidders tend to have similar feedback profiles. This technique looks for anomalies in the feedback data to determine whether an individual fits the profile of a shill. However, the problem with

this approach is that it is specific to eBay, and is based solely on the integrity of the feedback system.

One possible bidding strategy to combat shilling is bid sniping. Sniping refers to a bid that is submitted just before the auction terminates. The intent of the bid is to prevent other bidders from outbidding it by denying them time to react. Sniping cannot prevent shilling from occurring in an auction. However, it does protect a bidder from being shilled even further.

Existing commercial Auctioneers permit sniping, but discourage the practice. If every bidder engaged in sniping, the auction would essentially become a sealed bid auction. A sealed bid auction differs to an English auction in that no one knows the value of any one else's bid. In contrast, English auctions are open bid auctions as everyone knows everyone else's bid. Instead eBay recommends that bidders bid early using their proxy bidding system.

The proxy bidding system allows a bidder to specify a maximum amount they are willing to pay. If a rival enters a new high bid, the proxy bidding system will minimally outbid it on behalf of the bidder until the maximum amount is reached. However, if all bidders use proxy bidding, the auction becomes a second price auction. Second price auctions differ to English auction in that the winner only has to pay the price of the second highest bid (i.e., the highest losing bid).

We stated in Section II that many of the existing shill detection methods are not public. It can be argued that it is dangerous to give a shill bidder knowledge of the inner workings of the detection method as this might allow them to subvert the system. However, the security (or ability) of the system should not rest with keeping it secret. In cryptography it is usually the case that confidence in a proposed security technique is only acquired over time if it withstands public scrutiny and attacks against its security. Likewise, if an adversary is allowed to attack the shill detection methods, then the confidence in a scheme will grow and ultimately better detection techniques can evolve.

VI. THE SHILL SCORE

This section presents an algorithm to detect a shill by looking for common shill behaviour. Our algorithm considers the case with **one shill**. The goal of the algorithm is to determine which bidder is most likely to be the shill out of a group of ℓ bidders. A bidder is examined over m auctions held by the same seller for the behaviour outlined in Section IV. Each characteristic of shill behaviour is given a rating which is combined to form the bidder's shill score. The shill score gives a bidder a value between 0 and 10. The closer the shill score is to 10, the more likely that the bidder is a shill.

The remainder of this section describes how to quantify a shill's characteristics to deduce a shill score. However, to do this some preliminary notation is required:

Let $L = \{1, \dots, \ell\}$ be the set of bidder numbers;

$$|L| = \ell$$

Let $M = \{1, \dots, m\}$ be the set of auction numbers;

$$|M| = m$$

For $i \in L$, let $M^i = \{j \mid j \in M, \text{ bidder } i \text{ bids in auction } j\}$;

$$|M^i| = m^i$$

For $j \in M$, let $N_j = \{1, \dots, n_j\}$ be the bid numbers (e.g., 1st bid, 2nd bid, etc.) in auction j ;

$$|N_j| = n_j$$

For $i \in L, j \in M$, let $N_j^i = \{k \mid k \in N_j, \text{ bidder } i \text{ makes the } k\text{th bid in auction } j\}$;

$$|N_j^i| = n_j^i$$

For $i \in L$, let $W^i = \{j \mid j \in M, \text{ bidder } i \text{ wins auction } j\}$;

$$|W^i| = w^i$$

For $j \in M, k \in N_j$, let $t_{j,k}$ be the time of the k^{th} bid in auction j , and let

$$T_j = \{t_{j,k} \mid j \in M, k \in N_j\}$$

be the set of bid times for auction j .

For $i \in L, j \in M, k \in N_j$, let $t_{j,k}^i$ be the time of the k^{th} bid in the auction j by bidder i . Let

$$T_j^i = \{t_{j,k}^i \mid i \in L, j \in M, k \in N_j\}$$

be the set of bid times for bidder i in auction j ;

$$|T_j^i| = n_j^i$$

For $j \in M, k \in N_j$, let $p_{j,k}$ be the value of the k^{th} bid in auction j , and let

$$P_j = \{p_{j,k} \mid j \in M, k \in N_j\}$$

be the set of bid values in auction j ;

$$|P_j| = n_j$$

For $i \in L, j \in M, k \in N_j$, let $p_{j,k}^i$ be the value of the k^{th} bid in the auction j by bidder i . Let

$$P_j^i = \{p_{j,k}^i \mid i \in L, j \in M, k \in N_j\}$$

be the set of bid values for bidder i in auction j ;

$$|P_j^i| = n_j^i$$

1) α Rating: As stated in Section IV, a shill is likely to participate exclusively in auctions held by one particular seller. The first step in detecting shilling is to count the number of the seller's auctions bidder i has participated in. Any auctions that bidder i has won (denoted as w^i) are excluded.

Given a seller which has held m auctions, the percentage α^i , of auctions bidder i has participated in is calculated as:

$$\alpha^i = (m^i - w^i)/m$$

where $0 \leq \alpha^i \leq 1$.

In general, the α rating will be higher for a shill compared to a legitimate bidder.

2) β Rating: An aggressive shill tends to submit a lot of bids as they are continually trying to outbid others. This might involve the shill bidding as often as every second bid (i.e., after every legitimate bid). Table II illustrates this behaviour. In the worst case a shill must submit $\lfloor n_j/2 \rfloor$ bids, where n_j is the total number of bids for an auction. If a shill bids any more bids than this, they will win the auction.

This bound holds regardless of how many bidders there are in the auction. Table III shows an auction with two normal bidders and one shill bidder. At most the shill bidder can submit $\lfloor n_j/2 \rfloor$ bids. In the best case, a shill will not have to bid because the legitimate bidders have already bid up the price well beyond the expected payoff from shilling.

The β rating indicates the average percentage of bids that bidder i has submitted throughout all the auctions he/she has participated in. The percentage for an individual auction is calculated in terms of the worst case bound of $\lfloor n_j/2 \rfloor$ bids. If a bidder has won an auction, then his/her β rating for that auction is zero.

An individual's β rating is calculated as follows:

Count the number of auctions m^i , that bidder i has participated in. For all auctions $j \in M^i$, that bidder i has participated in, perform the following:

Let β_j^i be bidder i 's β rating for auction j .

- 1) If bidder i won auction j , then $M^i = M^i - j$ (and $m^i = m^i - 1$) and skip the remaining numbered steps. Otherwise proceed to step 2.
- 2) Calculate the number of bids, n_j , that have been submitted by all bidders in auction j .
- 3) Calculate the worst case number of bids ω_j , for auction j :

$$\omega_j = \lfloor n_j/2 \rfloor$$

- 4) Calculate the number of bids, n_j^i , bidder i has submitted in auction j .
- 5) Calculate the percentage of bids made by bidder i in terms of the worst case bound ω_j :

$$\beta_j^i = n_j^i/\omega_j$$

Finally, bidder i 's β rating is calculated as the average of his/her β_j^i ratings for all auctions m_i , participated in:

If $m^i = 0$, then

$$\beta^i = 0$$

Else

$$\beta^i = \frac{1}{m^i} \sum_{j \in M^i} \beta_j^i$$

where $0 \leq \beta^i \leq 1$.

In general, the β rating will be high for an aggressive shill.

TABLE II
AUCTION WITH TWO BIDDERS - ONE BIDDER (b_1), ONE SHILL (s_1)

n_j	b_1	s_1	bid sequence	b_1	s_1
3	2	1	b_1, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$
4	2	2	s_1, b_1, s_1, b_1	$n_j/2$	$n_j/2$
5	3	2	b_1, s_1, b_1, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$
6	3	3	$s_1, b_1, s_1, b_1, s_1, b_1$	$n_j/2$	$n_j/2$

TABLE III
AUCTION WITH THREE BIDDERS - TWO BIDDERS (b_1, b_2), ONE SHILL (s_1)

n_j	b_1	b_2	s_1	bid sequence	b_1	b_2	s_1
3	1	1	1	b_1, s_1, b_2	$\lceil n_j/2 \rceil$		$\lfloor n_j/2 \rfloor$
4	1	1	2	s_1, b_1, s_1, b_2	$n_j/2$		$n_j/2$
5	2	1	2	b_1, s_1, b_2, s_1, b_1	$\lceil n_j/2 \rceil$		$\lfloor n_j/2 \rfloor$
6	2	1	3	$s_1, b_1, s_1, b_2, s_1, b_1$	$n_j/2$		$n_j/2$

3) γ Rating: A shill's secondary goal is to avoid winning, as the auction will have to be repeated (with cost to the seller). To measure this, the simplest approach is to determine how many times bidder i has actually won over the auctions they have participated in.

For a bidder i that has won w^i auctions, the percentage γ , of wins is calculated as follows:

$$\gamma^i = 1 - (w^i/m^i)$$

where $0 \leq \gamma^i \leq 1$.

In general a shill will have a high γ rating. However, calculating the γ rating in this manner is biased towards bidders who only bid one or two bids and bidders that have participated in a small number of auctions.

We refer to a bidder who has only submitted one bid as a *one-time* bidder. At present, a one-time bidder that loses the auction attains a γ rating of one. This artificially inflates the bidder's shill score even though he/she has only bid once.

Our own experiments using commercial auction data have shown that many of the bidders participating in a set of auctions by a particular seller, are often one-time bidders. Calculating the γ rating in this manner gives misleading results when large numbers of one-time bidders are present.

To rectify this problem, the γ rating must also consider the number of losses, l^i , incurred by bidder i , where $l^i = m^i - w^i$. As it is common for a bidder to lose auctions it makes no sense to penalise small numbers of losses. Instead, suspicion is only raised when the number of losses becomes large in comparison to the number of winnings. One-time bidders can be identified as bidders with no winnings and only one loss. Such bidders are given a γ rating of zero.

The γ rating is altered to permit a 20% loss to winnings ratio without penalty. This allows bidders with a small number of losses to be ignored. The γ rating is further altered to ensure it is scalable depending on the number of auctions

a bidder participates in. As the number of auctions that bidder i participates in increases, the penalties for losing an auction increase. This effectively identifies bidders who have participated in an excessive number of auctions and have little or no winnings.

The γ rating is calculated as follows:

$$\gamma^i = 1 - \left(\frac{5 \times (w^i + 0.2)}{l^i} \right)$$

where $\gamma^i < 1$.

In the case where $\gamma^i < 0$, then $\gamma^i = 0$. Thus $0 \leq \gamma^i < 1$.

4) δ Rating: A shill wants to give other bidders as much time as possible to consider a shill bid. Therefore a shill generally submits a new bid within a small time period of a rival bid.

This behaviour can be measured by observing inter bid times for all bidders. The average inter bid time is found for bidder i across the auctions they have participated in. Bidders who wait longer between bids have a lower average inter bid time score.

For all auctions $j \in M^i$, that bidder i has participated in, perform the following:

Let δ_j^i be bidder i 's δ rating for auction j .

- 1) If bidder i won auction j , then $M^i = M^i - j$ (and $m^i = m^i - 1$) and skip the remaining numbered steps. Otherwise proceed to step 2.
- 2) Calculate the inter bid time for each bid, k , submitted by bidder i in auction j :

$$\Delta t_{j,k}^i = \begin{cases} 0 & k = 1 \\ t_{j,k}^i - t_{j,k-1} & k > 1, k \in N_j \end{cases}$$

where $t_{j,k-1}$ is the time of a previous bid (by a rival bidder).

- 3) Calculate the average inter bid time for bidder i :

$$\bar{\Delta t}_j^i = \frac{1}{n_j^i} \sum_{k \in N_j^i} \Delta t_{j,k}^i$$

- 4) The maximum and minimum average inter bid times for auction j are:

$$\begin{aligned} \bar{\Delta t}_j^{max} &= \max_{i \in L} \{\bar{\Delta t}_j^i\} \\ \bar{\Delta t}_j^{min} &= \min_{i \in L} \{\bar{\Delta t}_j^i\} \end{aligned}$$

Normalise to find bidder i 's δ_j^i rating for auction j

$$\delta_j^i = \frac{\bar{\Delta t}_j^i - \bar{\Delta t}_j^{min}}{\bar{\Delta t}_j^{max} - \bar{\Delta t}_j^{min}}$$

Finally, calculate the average of the normalised inter bid times for bidder i over all auctions participated in:

$$\delta^i = 1 - \left(\frac{1}{m^i} \sum_{j \in M^i} \delta_j^i \right)$$

where $0 \leq \delta^i \leq 1$.

In general, the δ rating will be higher for a skill in comparison to a legitimate bidder.

5) ϵ Rating: A skill that outbids by a large amount increases the risk of losing the auction. Therefore a skill tends to only bid the minimal amount to stay ahead of the leading bid.

This behaviour can be measured by observing inter bid increments for all bidders. The average inter bid increment is found for bidder i across the auctions they have participated in. Bidders who submit smaller bid increments have a lower average inter bid increment score.

For all auctions $j \in M^i$, that bidder i has participated in, perform the following:

Let ϵ_j^i be bidder i 's ϵ rating for auction j .

- 1) If bidder i won auction j , then $M^i = M^i - j$ (and $m^i = m^i - 1$) and skip the remaining numbered steps. Otherwise proceed to step 2.
- 2) Calculate the inter bid increment for each bid, k , submitted by bidder i in auction j :

$$\Delta p_{j,k}^i = \begin{cases} 0 & k = 1 \\ p_{j,k}^i - p_{j,k-1} & k > 1, k \in N_j \end{cases}$$

where $p_{j,k-1}$ is the price of a previous bid (by a rival bidder).

- 3) Calculate the average inter bid increment for bidder i :

$$\bar{\Delta p}_j^i = \frac{1}{n_j^i} \sum_{k \in N_j^i} \Delta p_{j,k}^i$$

- 4) The average maximum and minimum inter bid increments for auction j are:

$$\begin{aligned} \bar{\Delta p}_j^{max} &= \max_{i \in L} \{\bar{\Delta p}_j^i\} \\ \bar{\Delta p}_j^{min} &= \min_{i \in L} \{\bar{\Delta p}_j^i\} \end{aligned}$$

Normalise to find bidder i 's ϵ_j^i rating for auction j :

$$\epsilon_j^i = \frac{\bar{\Delta p}_j^i - \bar{\Delta p}_j^{min}}{\bar{\Delta p}_j^{max} - \bar{\Delta p}_j^{min}}$$

Finally, calculate the average of the normalised inter bid increments for bidder i over all auctions participated in:

$$\epsilon^i = 1 - \left(\frac{1}{m^i} \sum_{j \in M^i} \epsilon_j^i \right)$$

where $0 \leq \epsilon^i \leq 1$.

In general, the ϵ rating will be higher for a skill than that of a legitimate bidder.

6) ζ Rating: If a skill bids late in an auction, he/she risks losing the auction by not being outbid in time. Therefore it is in the skill's best interests to bid early in an auction as it is less risky and maximises the amount of influence the skill has over an auction.

The ζ rating indicates how early in an auction bidder i commenced bidding. This is calculated as the difference between the auction's expiration time and the time that bidder i first submitted a bid. This result is normalised against all other participants in the auction and averaged over all the auctions that bidder i participated in.

An individual's ζ rating is calculated as follows:

Count the number of auctions m^i , that bidder i has participated in. For all auctions $j \in M^i$, that bidder i has participated in, perform the following:

Let ζ_j^i be bidder i 's ζ rating for auction j .

- 1) If bidder i won auction j , then $M^i = M^i - j$ (and $m^i = m^i - 1$) and skip the remaining numbered steps. Otherwise proceed to step 2.
- 2) Calculate the difference between auction j 's expiration time, t_j , and the time, t_{j,k_0}^i of the first bid submitted by bidder i in auction j .

$$\Delta t_j^i = t_j - t_{j,k_0}^i$$

- 3) The maximum and minimum average final bid time differences for auction j are:

$$\begin{aligned} \Delta t_j^{max} &= \max_{i \in L} \{\Delta t_j^i\} \\ \Delta t_j^{min} &= \min_{i \in L} \{\Delta t_j^i\} \end{aligned}$$

Normalise to find bidder i 's ζ_j^i rating for auction j

$$\zeta_j^i = \frac{\Delta t_j^i - \Delta t_j^{min}}{\Delta t_j^{max} - \Delta t_j^{min}}$$

Finally, calculate the average of the final bid time differences for bidder i over all auctions participated in:

$$\zeta^i = \frac{1}{m^i} \sum_{j \in M^i} \zeta_j^i$$

where $0 \leq \zeta^i \leq 1$.

In general, the ζ rating will be higher for a skill.

TABLE IV
AN EXAMPLE AUCTION WITH PROXY BIDDING

Bid #	b_{id}	Price	Time
8	b_5	\$ 17	3:01
7	b_4	\$ 16	2:54
2	b_2	\$ 15	0:33
6	b_3	\$ 9	2:00
5	b_3	\$ 7	1:55
4	b_3	\$ 5	1:47
3	b_3	\$ 3	1:40
1	b_1	\$ 1	0:05

7) *Pre-conditioning the Data*: Commercial auctions such as those on eBay do not adhere to the English auction format used in this paper. Before executing the shill detection algorithm, the data from commercial auctions must be ‘pre-conditioned’. Pre-conditioning the data refers to the process of getting the data into the correct format for the algorithm. This also involves modifying the algorithm to deal with proxy bidding.

As stated earlier, a proxy bid allows a user to enter a maximum bid that they are willing to pay. The proxy bidding system will then marginally outbid any new high bid up until the user’s maximum value. We refer to the first proxy bid submitted by the bidder as the *initial proxy bid*. A bid placed by the proxy system is referred to as a *proxy generated bid*. If the winning bid is less than the bidder’s maximum value, then the bidder gets the item for the lower value. Once a proxy bid is made, it cannot be retracted.

Table IV shows an example auction using proxy bidding. In this example b_2 is the proxy bidder. The bid number indicates the order the bids were received. The auction history does not display the proxy generated bids. The initial proxy bid is only revealed when it is outbid. If a proxy generated bid wins the auction, its value is shown, but the initial proxy bid is not revealed. If two proxy bidders are competing, the lower of the two bids is listed and the proxy generated bid of the higher valued bid is listed.

A shill is unlikely to use proxy bidding. This is the same justification behind the ϵ rating which states that a shill is more likely to bid in smaller increments rather than one large up-front bid. Proxy bidding encourages bidders to bid large amounts up-front. However, this behaviour is risky for a shill.

A safer approach is for a shill to edge up a proxy bid until it reaches its maximum. If a shill can predict what the value of the initial proxy bid is, then the shill can seize the full value of the proxy bid.

When a shill encounters a proxy bid, this presents the shill with an accessible target to shill. Upon outbidding a proxy generated bid, the shill is provided with instant feedback. That is, the shill is either outbid immediately by a new proxy generated bid or the shill has outbid the proxy bid. Such instant feedback allows a shill to inflate the price quickly and in a controlled incremental manner. Upon submitting a proxy bid, the bidder is powerless to prevent a shill from edging the price

TABLE V
EXAMPLE AUCTION RE-ORDERED ACCORDING TO THE TIME SUBMITTED

Bid #	b_{id}	Price	Time
8	b_5	\$ 17	3:01
7	b_4	\$ 16	2:54
6	b_3	\$ 9	2:00
5	b_3	\$ 7	1:55
4	b_3	\$ 5	1:47
3	b_3	\$ 3	1:40
2	b_2	\$ 15	0:33
1	b_1	\$ 1	0:05

up as the initial bid cannot be retracted.

Note that alternately a shill might use a proxy bid up-front in order to inflate the price to a given value. This can be likened to an undisclosed reserve under the guise of a legitimate bid. Such behaviour is essentially benign shilling. If a shill uses multiple proxy bids throughout an auction, they must do so sparingly. Once again it is risky for the shill to set an initial proxy bid too high. This effectively limits the usefulness of proxy bidding to a shill.

To detect shill bidders in commercial online auctions, the algorithm must be altered to deal with proxy bidders. The first step to achieving this goal is to re-order the auction data based on time rather than bid value. This facilitates the identification of initial proxy bids.

An initial proxy bid can be identified as a bid that has a higher value than a subsequently placed bid. The re-ordered data from the example auction is shown in Table V. Bid number 2 submitted by b_2 clearly stands out as an initial proxy bid. This bid is for \$15 whereas the subsequent bid (i.e., bid number 3) is lower at \$3.

The total number of generated proxy bids (not shown in the bidding history) can be determined by counting the number of subsequent bids that have a lower value than the initial proxy bid. In Table V there are four rival bids after b_2 ’s initial proxy bid. Therefore there are also four proxy generated bids.

To ensure that a proxy bidder’s β rating does not exceed one, the number of proxy generated bids for auction j is added to the total number of bids n_j . When a proxy bid is used, the β rating only penalises the proxy bidder for one bid (i.e., the initial proxy bid). Proxy generated bids are not included in the proxy bidder’s β rating. This functionality rewards and encourages proxy bidding.

A proxy bidder’s δ rating is only calculated for the initial proxy bid. Proxy generated bids are not included in the proxy bidder’s δ rating. This is because the timings for the proxy generated bid will always be zero as they are placed immediately once a rival bid of lesser value is received. This wrongly gives the appearance that proxy generated bids are being submitted by a shill bidder.

To allow a bidder’s ϵ rating to be determined, the values of the proxy generated bids must be calculated. As proxy generated bids are not included in the auction history, the shill

TABLE VI
 BID INCREMENTS

Current Price	Bid Increment
\$0.00 - \$0.99	\$0.05
\$1.00 - \$4.99	\$0.25
\$5.00 - \$24.99	\$0.50
\$25.00 - \$99.99	\$1.00
\$100.00 - \$249.99	\$2.50
\$250.00 - \$499.99	\$5.00
\$500.00 - \$999.99	\$10.00
\$1000.00 - \$2499.99	\$25.00
\$2500.00 - \$4999.99	\$50.00
\$5000.00 and up	\$100.00

score algorithm cannot accurately determine the ϵ ratings for bidders in competition with a proxy bid. eBay has guidelines regarding the minimal amount required to outbid the current high bid. This is shown in Table VI as an increasing scale that is a function of the price (roughly 5%). These guidelines can be used to reproduce the missing proxy generated bid values.

The ϵ rating for the initial proxy bid is calculated as normal according to time. As the proxy bidder has bid a large amount up-front, this is manifested in his/her ϵ rating. Proxy generated bids are not included in the proxy bidder's ϵ rating.

The bid immediately following the initial bid is calculated in terms of the proxy generated bid produced. When an initial bid is placed, it is only listed as the minimum amount required to win. That is, the first proxy generated bid is displayed. The rival bidders ϵ value is calculated in terms of the first proxy generated bid rather than the initial bid. All subsequent rival bids are calculated in terms of the proxy generated bids they have outbid.

8) *Calculating the Skill Score:* The following algorithm is run on a seller's auction history for each bidder that participated in any of the seller's auctions:

- 1) Pre-condition the data.
- 2) Calculate α rating.
- 3) Calculate β rating.
- 4) Calculate γ rating.
- 5) Calculate δ rating.
- 6) Calculate ϵ rating.
- 7) Calculate ζ rating.
- 8) Calculate skill score.

Each bidder obtains a skill score between 0 and 10, representing their likelihood of being a skill. This is calculated as follows:

$$score = \frac{\theta_1\alpha + \theta_2\beta + \theta_3\gamma + \theta_4\delta + \theta_5\epsilon + \theta_6\zeta}{\theta_1 + \theta_2 + \theta_3 + \theta_4 + \theta_5 + \theta_6} \times 10$$

where θ is the weighting associated with each rating. The bidder with the highest skill score is considered most likely to be the skill.

The weights used in this paper are: $\theta_1 = 9$, $\theta_2 = 2$, $\theta_3 = 5$, $\theta_4 = 2$, $\theta_5 = 2$ and $\theta_6 = 2$. These weightings were obtained

by experimenting with simulated auction data and using the following justifications:

A bidder's α rating is given the highest weighting. This is done to eliminate one-time bidders. In the instance where a one-time bidder has high δ and ϵ ratings, the bidder will be allocated a high skill score despite clearly not fitting the profile of a skill.

A bidder's γ rating receives the next highest weighting. When a bidder has a high α rating and a low γ rating, the bidder is clearly bidding to win. This is irrespective of the bidder participating in so many auctions. Employing this strategy, a skill would not be very successful.

Once a bidder has initially been detected by his/her α and γ ratings, the next significant factors are the β , δ , ϵ and ζ ratings. These are used to see if the bidding behaviour matches the regular behaviour of a skill.

The β rating measures the number of bids submitted in terms of $\lfloor n_j/2 \rfloor$. If a skill is aggressive, β will be high. However, if the skill is benign, β will be low. As a result, the β rating can present mixed results depending on the type of skill behaviour employed. For this reason, the β rating is given a lower weighting than the α and γ ratings.

The δ and ϵ weightings are also lower than the α and γ weightings due to the effect of one-time bidders. That is, if a bidder only bids once, placing the bid quickly after the current highest bid and by the minimal amount required, then the bidder will have high δ and ϵ ratings. As the bidder does not bid again, his/her δ and ϵ ratings will always remain high and not average down if the bidder were to later submit slower, larger bids. Therefore, weighting δ and ϵ highly results in many one-time bidders scoring high overall even though such bidders are clearly not skills (evidenced by a low α rating).

The ζ rating is also weighted at two as it is not more influential than the other bidding characteristic ratings. Instead all the bidding characteristic weightings must be examined as a group to determine if the bidder's bidding behaviour fits the profile of a skill.

Note that the weightings must be adjusted according to the number of auctions held. For small values of m the α and γ ratings are less effective. For example, when $m \leq 3$ the α rating for all bidders will be high and the γ rating will be zero for all of the bidders. This distorts a bidder's skill score and as a result cannot be relied upon. When $m \leq 3$ the weightings for the α and γ ratings are zero.

In the case where $m \leq 3$, the bidding characteristic ratings are more heavily relied on to indicate the level of skill-like behaviour a bidder possesses.

VII. PERFORMANCE

This section describes how the algorithm performed on real auction data. Two types of data were obtained. The first was from a series of simulated auction trials. The second was from commercial online auctions. Each type will be explained in turn.

TABLE VII
SIMULATED AUCTION WITH ONE SHILL

Bid #	b_{id}	Price	Time
25	Ness	\$ 91.00	2005-08-11 15:58:35
24	wayne	\$ 90.00	2005-08-11 15:44:04
23	townsville	\$ 80.00	2005-08-11 15:22:46
22	Ness	\$ 56.00	2005-08-11 15:16:16
21	townsville	\$ 55.00	2005-08-11 15:07:28
20	Shelly	\$ 50.50	2005-08-11 14:55:34
19	townsville	\$ 50.00	2005-08-11 14:54:46
18	Shelly	\$ 39.39	2005-08-11 14:12:55
17	buzzcook	\$ 39.00	2005-08-11 14:11:02
16	Shelly	\$ 38.38	2005-08-11 14:04:43
15	marianne	\$ 38.00	2005-08-11 13:49:17
14	Shelly	\$ 36.72	2005-08-11 13:40:26
13	Soraya	\$ 36.36	2005-08-11 13:34:23
12	wayne	\$ 36.00	2005-08-11 08:31:09
11	Shelly	\$ 35.35	2005-08-10 14:47:40
10	Soraya	\$ 35.00	2005-08-10 14:41:33
9	Shelly	\$ 33.33	2005-08-10 14:34:15
8	wayne	\$ 33.00	2005-08-10 14:25:49
7	Ness	\$ 32.00	2005-08-10 14:20:13
6	Shelly	\$ 31.31	2005-08-10 11:11:21
5	wayne	\$ 31.00	2005-08-10 11:09:55
4	Soraya	\$ 30.50	2005-08-10 08:25:17
3	Ness	\$ 30.00	2005-08-09 16:42:48
2	marianne	\$ 0.25	2005-08-09 16:25:13
1	Shelly	\$ 0.20	2005-08-09 16:01:18

A. Simulated Auctions

A series of simulated auctions were conducted to obtain auction data. Thirty-nine auctions were performed on an auction server which emulates existing commercial online auction websites. Details of the software used can be found in [Trevathan and Read 2006], and the auction server itself is available at *auction.maths.jcu.edu.au*. Each auction was for a different item and all auctions were considered to be for one seller (consistent with the algorithm)

The auction proceedings involved twenty-six people. Each bidder was given a random amount of fake money at the beginning of the trials. A bidder was free to bid his/her valuation in any auction provided that collectively the amount bid did not exceed the initial amount provided. When a bidder won an auction, the balance of his/her account was reduced by the value of the winning bid. A bidder's goal was to win while also trying to save his/her money.

The shill's goal was to force a bidder into spending as much of his/her money as possible. Bidders were not informed that shilling was occurring. Furthermore, the shill had no knowledge of how much money bidders had.

Table VII shows the proceedings of one of the simulated auctions. The item auctioned was Dire Straits' Greatest Hits DVD. This consisted of a DVD of the band's video clips and two CDs. The retail value was approximately \$50 AUD. The shill bidder is 'Shelly'.

It quickly became apparent that bidders took the auctions seriously. The simulation captured the essence of a real auction where the competitive desire to win induced bidders to

TABLE VIII
SHILL SCORES FOR SIMULATED AUCTIONS WITH ONE SHILL BIDDER

Bidder	α	β	γ	δ	ϵ	ζ	Shill Score
Shelly	1.00	0.70	0.90	0.96	0.88	0.79	9.16
townsville	0.60	0.21	0.83	0.86	0.65	0.60	6.46
Marie	0.60	0.11	0.83	0.45	0.89	0.20	5.84
jc112425	0.40	0.20	0.75	0.75	0.72	0.79	5.58
wayne	0.20	0.42	0.50	0.75	0.93	0.66	4.46
Ness	0.40	0.48	0.00	0.88	0.90	0.85	4.46
marianne	0.60	0.18	0.00	0.62	0.91	0.39	4.36
Soraya	0.40	0.32	0.00	0.91	0.85	0.78	4.25
Bear	0.20	0.22	0.50	0.98	0.16	0.43	3.58
brenda	0.20	0.59	0.00	0.97	0.97	0.41	3.49
joe	0.10	0.67	0.00	0.62	0.92	1.00	3.32
banana	0.10	0.11	0.00	1.00	1.00	1.00	3.24
ronisanidiot	0.10	0.10	0.00	0.95	0.98	0.76	2.95
phrdw	0.10	0.10	0.00	0.99	0.99	0.68	2.92
buzzcook	0.10	0.08	0.00	1.00	0.99	0.01	2.30
groper	0.00	0.00	0.00	0.00	0.00	0.00	0.00

spend more money than was intended. Bidders that won an auction were elated, even though they did not obtain anything tangible. Other bidders such as 'brenda' vowed that they were determined to win all of the auctions.

The simulation showed that while the shill's strategy was to marginally increase the bid, legitimate bidders tended to increase the bid by larger values. Furthermore, legitimate bidders who were focused on winning an item regardless of the cost tended to bid large amounts up front. In addition, many bidders engaged in rallies amongst themselves and with the shill bidder. During such rallies legitimate bidders tended to bid in larger increments than the shill's marginal increases.

For the auction shown in Table VII, the shill (i.e., 'Shelly') pushed the price up to \$50.50. At this stage the auction was one hour from terminating. Although bidding was constant, the reserve price had been met therefore it was in the shill's interests to discontinue. The winning bid for the item was \$91.

There were three types of tests performed: with one shill, without shilling and with benign shilling. The shill detection algorithm was run on all of the tests to determine its effectiveness in catching shills and the likelihood of it incriminating innocent bidders.

The first test involved ten auctions and one shill bidder. Table VIII shows the shill scores for each bidder. The shill (i.e., 'Shelly') is clearly identified as the bidder that has engaged in the most shill-like behaviour. The shill was the only bidder to score consistently on the behavioural characteristics that optimally inflate the price for the seller. Figure 1 A gives a graphical representation of the shill scores.

The second test also involved ten auctions. However, unlike the previous test, no intentional shilling behaviour was engaged in. Shill scores for each bidder are shown in Table IX and graphically in Figure 1 B. The purpose of the test was to gauge regular bidding behaviour.

At some stages during an auction, a legitimate bidder's

TABLE IX
SHILL SCORES FOR SIMULATED AUCTIONS WITHOUT SHILLING

Bidder	α	β	γ	δ	ϵ	ζ	Shill Score
jc112425	0.40	0.36	0.75	0.99	0.90	0.99	6.28
Ness	0.30	0.64	0.67	0.98	0.87	1.00	5.92
banana	0.20	0.75	0.50	1.00	1.00	1.00	5.36
dee	0.20	0.42	0.50	0.97	0.91	0.43	4.44
joe	0.10	1.00	0.00	1.00	1.00	1.00	4.05
brenda	0.30	0.38	0.00	0.84	0.85	0.17	3.27
jc149722	0.20	0.25	0.50	0.95	0.00	0.22	3.25
marianne	0.20	0.30	0.00	0.81	0.83	0.73	3.24
Bear	0.20	0.42	0.00	0.98	0.76	0.49	3.23
Sarah	0.10	1.00	0.00	0.96	0.68	0.46	3.22
wayne	0.10	0.50	0.00	0.98	0.63	0.95	3.18
townsville	0.30	0.44	0.00	0.50	0.50	0.28	2.80
Soraya	0.10	0.50	0.00	0.62	0.79	0.50	2.60
jai	0.10	0.17	0.00	0.99	0.93	0.00	2.31
Marie	0.10	0.14	0.00	0.95	0.92	0.07	2.30
buzzcook	0.10	0.14	0.00	0.98	0.87	0.01	2.23
Jarrold	0.10	1.00	0.00	0.00	0.00	1.00	2.23
metoo	0.10	0.17	0.00	0.89	0.93	0.00	2.22

TABLE X
SHILL SCORES FOR SIMULATED AUCTIONS WITH BENIGN SHILLING

Bidder	α	β	γ	δ	ϵ	ζ	Shill Score
joe	1.00	0.68	0.95	0.99	1.00	1.00	9.58
Marie	0.37	0.45	0.86	0.62	0.59	0.66	5.56
brenda	0.42	0.52	0.25	0.65	0.70	0.59	4.52
alan	0.11	0.25	0.50	0.85	0.95	0.84	4.20
ronisanidiot	0.16	0.24	0.67	0.97	0.63	0.35	4.15
banana	0.11	1.00	0.50	0.89	0.02	0.90	4.13
Mezza	0.16	0.26	0.67	0.82	0.94	0.07	4.07
Jarrold	0.16	0.56	0.67	0.75	0.33	0.43	4.04
jc112425	0.26	0.28	0.00	0.61	0.81	0.60	3.16
Shelly	0.05	0.25	0.00	0.86	0.98	0.90	2.93
Soraya	0.16	0.39	0.00	0.96	0.89	0.17	2.84
Sarah	0.16	0.31	0.00	0.65	0.91	0.52	2.82
wayne	0.11	0.29	0.00	0.99	0.95	0.05	2.50
Davey	0.05	0.13	0.00	0.99	1.00	0.10	2.22
phrdw	0.11	0.38	0.50	0.00	0.00	0.25	2.14
Bear	0.21	0.24	0.00	0.53	0.51	0.08	2.09
Kain0_0	0.05	0.13	0.00	0.78	0.00	0.12	1.14
groper	0.00	0.00	0.00	0.00	0.00	0.00	0.00

behaviour can emulate that of a shill. This raises the possibility of classifying innocent behaviour as shilling. During the trial some bidders participated in many of the auctions and also engaged in rallies which pushed up the bid price (e.g., ‘brenda’, ‘marianne’, ‘Ness’, ‘townsville’ and ‘Soraya’). This behaviour effectively raised their shill scores. However, other factors such as number of winnings, large bid increments and slower bid times indicated that these bidders were bidding to win.

During both the first and second tests, one of the bidders (‘banana’ in this case) was tasked with introducing some initial bids into auctions where there had been no bidding. The idea was to try and stimulate bidding (i.e., benign shilling).

When a bidder is the first to bid in an auction, there is no

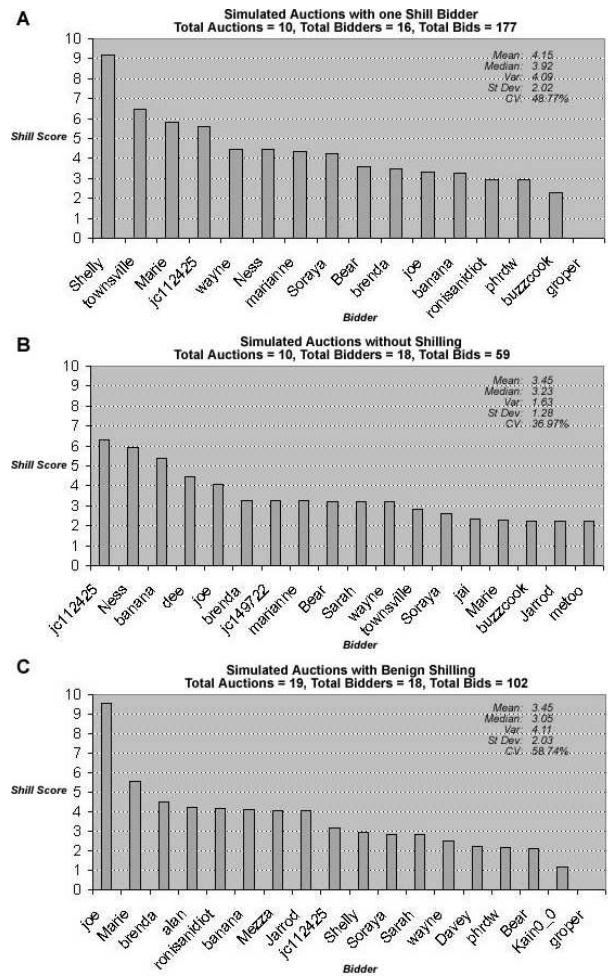


Fig. 1. Shill Scores for Simulated Auction Data

previous bid. As a result there is no inter bid time interval nor bid increment for the bid. The shill detection algorithm gives bids of this sort a δ and ϵ rating of one. If a benign shill continually bids first in an auction (to try ‘kick off’ the bidding) he/she will have high overall δ and ϵ ratings. Such a bidder will also have relatively high α and γ ratings if the behaviour continues over a series of auctions.

‘banana’ has one of the highest δ and ϵ ratings out of all the bidders in both tests. Since ‘banana’ only shilled in relatively few auctions, his α rating is low. A benign shill’s γ rating will always be high. In general, a benign shill will have a low β rating, except in situations where the auctions participated in have a low bid volume.

A third series of auction tests were conducted to further investigate benign shilling. This involved nineteen auctions, fifteen of which contained instances of benign shilling behaviour. Shill scores for each bidder are shown in Table X and graphically in Figure 1 C.

In this instance ‘joe’ is the shill bidder. ‘joe’ rates highly across all measures. The volume of bidding was low in several of the auctions ‘joe’ participated in. This accounts for why ‘joe’s’ β rating is reasonably high.

Figure 2 shows the average inter bid times and increments

TABLE XI
TOP FIVE SHILL SCORES FOR SELLER: SAVEKING

Bidder	α	β	γ	δ	ϵ	ζ	Shill Score
camotor1@aol.com	0.80	0.29	0.88	0.77	0.70	0.32	7.16
bacardi91	0.40	0.26	0.75	0.68	0.91	0.87	6.15
ypolina	0.30	0.34	0.67	0.55	0.59	0.18	4.25
janicerunning	0.20	0.24	0.50	0.96	0.75	0.15	3.86
s13g	0.20	0.12	0.50	0.86	0.89	0.11	3.76




Fig. 2. Time Intervals and Bid Increments for the Dire Straits' DVD Auction

for each bidder in the auction from Dire Straits' Greatest Hits DVD in Table VII. The shill bid the most during this auction and had a consistently low inter bid time and increment. Legitimate bidders that had inter bid times and increments less than the shill had submitted significantly fewer bids. When compared to the shill scores (see Table VIII) only bidders 'Shelly' and 'marianne' rank highly and warrant further investigation. While Figure 2 indicates which bidders are potential shills, it is insufficient to only look at an individual auction. For example, 'buzzcook' also rates poorly, but when his behaviour from several auctions is observed he has a low ranking shill score.

It is interesting to note that there was a significant degree of sniping occurring during all three tests. Many bidders such as 'brenda', 'buzzcook', 'Davey' and 'wayne' often employed this strategy to ensure that they won. It became evident that online auctions which terminate exclusively with an expiration time promote this sort of bidding. As this behaviour is totally the opposite of shilling, these bidders tended to have low shill scores.

These simulated tests are currently being repeated using real money. Furthermore, software bidding agents are being developed in order to conduct further simulations. A software bidding agent is a program that bids on behalf of a bidder according to a predefined strategy. This allows the shill detection algorithm to be tested on extremely large datasets.

B. Commercial Auctions

The algorithm was tested on commercial auction data. This refers to actual data collected from real auctions that were

performed by commercial online auctioneers such as eBay.

There is limited commercial data on shilling. Several commercial online auction sites were approached for auction data (i.e., bidding histories). We also stated in the request that we are "interested in auctions where it has either been confirmed or suspected that shilling has occurred."

Most commercial sites were reluctant to give out information for privacy reasons. Shilling is cognate to other types of fraud (e.g., credit card fraud). Companies are reluctant to admit shilling has occurred as it would tarnish their reputation. This can also partially justify why auction companies do not give out extensive bidding history records, especially those containing shilling.

The shill score algorithm was tested on commercial auction data obtained used in previous auction research by Jank and Shmueli [Jank and Shmueli 2004]. The auction data was for 150 auctions conducted on eBay for new Palm M515 PDAs. For a full description of how the data was gathered see [Jank and Shmueli 2004].

The commercial data differed to the simulated data in various ways. Firstly, the commercial data was for a single (but distinct) item (i.e., PDAs), whereas the simulated data were for different items. Second, there were a larger number of bidders in the commercial data. Thirdly, for the simulated data it was more common for a bidder to bid multiple times in a single auction. Finally, with the commercial data, a large number of users only bid once during an auction.

The data was analysed using the methods presented in this paper. Unlike the simulated data, it was unknown whether shilling had actually occurred. Therefore the algorithm determined whether shill behaviour was present solely by examining the bidding characteristics of the auction participants. This essentially provided the algorithm with a realistic test (i.e., the manner the algorithm would run in practice).

The remainder of this section gives the results for three different sellers of the commercial data. The first seller (named "saveking") conducted ten auctions involving seventy bidders. The top five shill scores for this seller are listed in Table XI and is shown graphically in Figure 3 A.

The bidder named 'camotor1@aol.com' has the highest score and warranted further investigation. On closer inspection, this bidder was found to have participated in 80% of the seller's auctions, and had not won any. The other bidders who scored high had also participated in quite a few auctions without winning, and had bid often and quickly in small

TABLE XII
TOP FIVE SHILL SCORES FOR SELLER: MICHAEL-33

Bidder	α	β	γ	δ	ϵ	ζ	Shill Score
chimam	0.53	0.26	0.88	0.82	0.52	0.95	6.50
rexth	0.27	0.34	0.75	0.65	0.85	0.74	5.07
bonnieellarae	0.27	0.21	0.75	0.65	0.83	0.50	4.79
charly.0	0.20	0.21	0.67	0.65	0.79	0.84	4.60
steve_howard	0.20	0.29	0.67	0.73	0.31	0.66	4.14

TABLE XIII
TOP FIVE SHILL SCORES FOR SELLER: SYSCHANNEL

Bidder	α	β	γ	δ	ϵ	ζ	Shill Score
kc10	0.63	0.21	0.80	0.97	0.86	0.97	7.11
bigroney1	0.50	0.26	0.75	0.87	0.63	0.21	5.54
reharrell	0.50	0.24	0.75	0.84	0.72	0.08	5.46
julhale	0.25	0.23	0.50	0.93	0.92	0.01	4.06
havokzprodigy	0.25	0.23	0.50	0.43	0.80	0.07	3.57

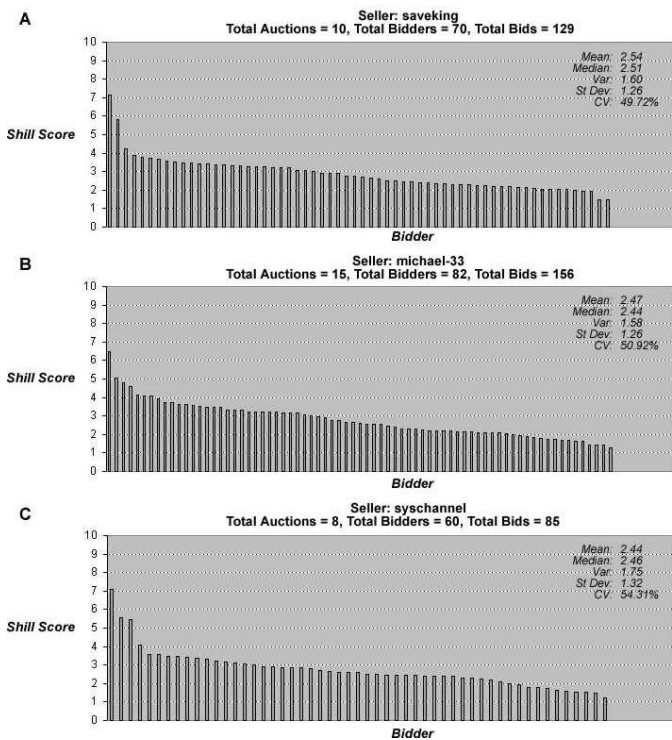


Fig. 3. Shill Scores for Commercial Auction Data

increments.

In contrast, there are bidders with shill scores of zero. These are bidders that won all of the auctions they participated in. Typical bidders had shill scores ranging between four and five.

The overall results for this seller suggests that there was a significant amount of price inflating bidding behaviour present. Out of seventy bidders, ‘camotor1@aol.com’ had the most suspicious behaviour. Given the extent of this behaviour compared to other bidders, this indicates that ‘camotor1@aol.com’ may

be engaging in aggressive shilling.

Table XII shows the top five shill scores for the seller named “michael-33”. The shill scores are calculated for 156 bidders over fifteen auctions. Shill scores for the bidders are shown graphically in Figure 3 B.

This series of auctions contained a lesser extent of price inflating bidding compared to the previous two sellers. The bidder with the highest shill score is ‘chimam’. This bidder has participated in 50% of the seller’s auctions and has not won any. However, this bidder does not have significantly high β , δ and ϵ ratings. Therefore it is questionable whether ‘chimam’ is actively engaging in shill behaviour. Instead, what the shill algorithm allows us to say is, “out of 156 bidders, this bidder engaged in the most significant price inflating behaviour”.

The other bidders listed in Table XII had mid range shill scores. On further inspection, these bidders had only participated in a small number of auctions and had the highest shill scores due to the “one-time” bidder phenomena described in Section VI.

Table XIII shows the top five shill scores for the seller named “syschannel”. The shill scores are calculated for fifty-nine bidders over eight auctions. This is shown graphically in Figure 3 C.

The bidder with the most suspicious behaviour is ‘kc10’. This bidder participated in the most auctions. Furthermore, ‘kc10’ (like ‘banana’ and ‘joe’ in the simulated auctions) was usually the first to bid in auctions that he/she participated in. This strongly indicates that ‘kc10’ may have been engaging in benign shilling behaviour.

The algorithm also singled out several other users for further investigation. For example, the bidders ‘bigroney1’ and ‘reharrell’ also exhibited dubious behaviour. This is evidenced by the large number of auctions participated in, lack of winnings and a larger percentage of bids submitted per auction than

other bidders. However, in comparison to ‘kc10’, the number of auctions, inter bid times and bid increments were slightly lower.

Bidders that had won any of the auctions had lowest skill scores. There were a large number of bidders that did not win, however, these bidders usually participated in only one auction. Such bidders tended to have a low range score (one to three) where their final ranking was heavily influenced by the inter bid times and bid increments.

The commercial data shows that the algorithm can effectively allow skill investigators to rule out many bidders and focus on those with the most significant skill-like behaviour.

The skill score algorithm is presently being tested on thousands of eBay auctions. We have written a program that parses an eBay HTML document extracting the auction data. This allows us to easily collect large amounts of commercial auction data via an automated process. We intend to continue collecting data in this manner to enable extensive testing of the skill score algorithm.

Security and anonymity is a significant issue for online auctions (see [Trevathan 2005], [Trevathan et al. 2005], [Trevathan et al. 2006]). Cryptographic mechanisms can be used to ensure that an auction is securely and fairly conducted. Furthermore, the identity of a bidder and the bid information can remain secret as long as the bidder doesn’t break the auction rules. In the case of a breach, a bidder’s identity can be revealed (with the equivalent of a court order) and appropriate recourse taken.

The skill detection algorithm presently requires all of this information to be available. The requirement for anonymity is evidenced by the responses from the commercial auction companies described earlier. At present, we are modifying the detection algorithm to operate in a secure and anonymous environment. The scheme allows the detection algorithm to be run without knowing the identities of bidders. However, when there is strong evidence to suggest shilling, a bidder’s identity can be revealed.

VIII. CONCLUSIONS

This paper presents the first published algorithm to detect the presence of skill bidding in online English auctions. Our algorithm targets core skill behaviour and calculates a score to indicate the likelihood that a bidder is a skill. To avoid detection a skill essentially must behave like a normal bidder. This renders a skill less effective, literally stopping skill bidding altogether.

The algorithm observes bidding patterns over a series of auctions held by a particular seller looking for skill-like behaviour. Bidding characteristics studied include how many auctions a bidder participated in, the number of bids submitted and the rate of auctions won. Furthermore, the algorithm inspects how quickly a bidder bids, how much a bidder outbids by and what stage in the auction a bidder commences bidding.

The ability of the skill score algorithm to effectively identify skill bidders was verified on simulated auctions. When faced with aggressive and benign shilling, the algorithm was able

to single out the skill bidder whilst exonerating legitimate bidders.

When tested on actual commercial data with no a priori knowledge of skills, the skill score algorithm detected definite skill behaviour. This was in the form of both aggressive and benign shilling. The algorithm also singled out other bidders with moderately suspicious behaviour.

The results of the tests showed that it is insufficient to rely on any single factor to determine whether a bidder is a skill. Instead a range of factors must be examined in relation to each other.

The skill score algorithm has significant practical and legal implications for commercial online auctions. This approach acts as both a skill detection mechanism and a deterrent to potential skills. Our algorithm serves as a basis for more sophisticated skill detection methods.

At present the algorithm only considers the case when there is one skill. However, in reality there may be multiple skills working in collusion. Multiple skills can employ sophisticated strategies which are more difficult to detect. Future work involves modifying the algorithm to better identify skill behaviour with colluding skills.

REFERENCES

- [Barbaro and Bracht 2005] Barbaro, S. and Bracht, B. 2005. Shilling, Squeezing, Sniping: Explaining late bidding in online second-price auctions. Working Paper.
- [Jank and Shmueli 2004] Jank and Shmueli. 2004. Dynamic Profiling of Online Auctions Using Curve Clustering. In *Proceedings of the Eleventh Annual Spring Research Conference (SRC) on Statistics in Industry and Technology*.
- [Kauffman and Wood 2003] Kauffman, R. and Wood, C. 2003. An Investigation of Premium Bidding in Online Auctions. Working Paper.
- [Kumar and Feldman 1998] Kumar, M. and Feldman, S. 1998. Internet Auctions. In *Proceedings of the Third USENIX Workshop on Electronic Commerce*. 49-60.
- [Schwartz and Dobrzynski 2002] Schwartz, J. and Dobrzynski, J. 2002. 3 men are charged with fraud in 1,100 art auctions on eBay. *The New York Times*.
- [Shah et al. 2002] Shah, H., Joshi, N. and Wurman, P. 2002. Mining for Bidding Strategies on eBay. In *SIGKDD’2002 Workshop on Web Mining for Usage Patterns and User Profiles*.
- [Stubblebine and Syverson 1999] Stubblebine, S. and Syverson, P. 1999. Fair On-line Auctions Without Special Trusted Parties. In *Proceedings of Financial Cryptography 1999*, (M. Franklin ed.), vol. 1648 of *Lecture Notes in Computer Science*. 230-240. Springer-Verlag.
- [Trevathan 2005] Trevathan, J. 2005. Security, Anonymity and Trust in Electronic Auctions. *ACM Crossroads*. Spring Edition, vol. 11.3, 3-9.
- [Trevathan et al. 2005] Trevathan, J., Ghodosi, H. and Read, W. 2005. Design Issues for Electronic Auctions. In *2nd International Conference on E-Business and Telecommunication Networks*. 340-347.
- [Trevathan et al. 2006] Trevathan, J., Ghodosi, H. and Read, W. 2006. An Anonymous and Secure Continuous Double Auction Scheme. In *39th International Hawaii Conference on System Sciences*. 125(1-12).
- [Trevathan and Read 2006] Trevathan, J., and Read, W. RAS: a system for supporting research in online auctions. *ACM Crossroads*, vol. 12.4, 23-30.
- [Wang et al. 2002] Wang, W., Hidvegi, Z. and Whinston, A. 2002. Skill Bidding in Multi-round Online Auctions. In *Proceedings of the 35th Hawaii International Conference on System Sciences*.
- [Wang et al. 2004] Wang, W., Hidvegi, Z. and Whinston, A. 2004. Skill-Proof Fee (SPF) Schedule: the Sunscreen against Seller Self-Collusion In Online English Auctions. Working Paper.

5.3 Detecting Collusive Shill Bidding

The shill score is designed for the instance where there is only one shill bidder (see Section 5.2). However, there are situations where there may be two or more shill bidders working in collusion with each other. Colluding shill bidders are able to engage in more sophisticated strategies that are potentially harder to detect. This paper proposes an extension to the shill score to detect colluding shill bidders. This is referred to as the *collusion graph*. The collusion graph either detects a colluding group, or forces the group members to act individually like a single shill, in which case they are detected by the shill score algorithm.

“Detecting Collusive Shill Bidding” [14] is to appear in the 4th International Conference on Information Technology - New Generations 2007 (ITNG). ITNG is an annual event focusing on state of the art technologies pertaining to digital information and communication. The applications of advanced information technology to such domains as astronomy, biology, education, geosciences, and health care are among topics of relevance to ITNG. Visionary ideas, theoretical and experimental results, as well as prototypes, designs, and tools that help the information readily flow to the user are of special interest. The conference features keynote speakers, the best student award, the PhD student forum and workshops/exhibits from industry, government and academia. Proceedings are published by the IEEE Computer Society. This paper is available online via citeseer⁴. ITNG 2007 is to be held in Las Vegas, U.S.A.

Publication in ITNG was made possible through funding by the Graduate Research Scholarship (GRS) and the School of Maths, Physics and IT.

⁴<http://citeseer.ist.psu.edu/>

Detecting Collusive Shill Bidding

Jarrod Trevathan

*School of Maths, Physics and IT
James Cook University*

Email: jarrod.trevathan@jcu.edu.au

Wayne Read

*School of Maths, Physics and IT
James Cook University*

Email: wayne.read@jcu.edu.au

Abstract—Shill bidding is where spurious bids are introduced into an auction to drive up the final price for the seller, thereby defrauding legitimate bidders. Trevathan and Read presented an algorithm to detect the presence of shill bidding in online auctions. The algorithm observes bidding patterns over a series of auctions, and gives each bidder a *shill score* to indicate the likelihood that they are engaging in shill behaviour. While the algorithm is able to accurately identify those with suspicious behaviour, it is designed for the instance where there is only one shill bidder. However, there are situations where there may be two or more shill bidders working in collusion with each other. Colluding shill bidders are able to engage in more sophisticated strategies that are harder to detect. This paper proposes a method for detecting colluding shill bidders, which is referred to as the collusion score. The collusion score, either detects a colluding group, or forces the colluders to act individually like a single shill, in which case they are detected by the shill score algorithm. The collusion score has been tested on simulated auction data and is able to successfully identify colluding shill bidders.

I. INTRODUCTION

Online auctions are a popular means to exchange items. However, to the unsuspecting, the auction process is fraught with peril. Auction fraud is an ever-increasing problem that is magnified in the online environment. For example, a seller might accept payment and not deliver an item, or misrepresents the item to be of greater value. Furthermore, auction participants might engage in practices such as siphoning, sniping, bid rigging or shilling, that are designed to influence the auction in a manner that disadvantages others (see [6]).

Shill bidding is the act of introducing fake bids into an auction on the seller's behalf, in order to artificially inflate an item's price. Bidders who engage in shilling are referred to as 'shills'. To win the item, a legitimate bidder must outbid a shill's price. If one of the shills accidentally wins, then the item is re-sold in a subsequent auction. Shill bidding is a problem as it forces legitimate bidders to pay significantly more.

The advent of online auctions such as eBay¹ and ubid² have made shill bidding much easier. This is due to bidders not being physically present, which allows a shill to anonymously influence the bidding process. Furthermore, it is relatively simple for a seller to register under many aliases, and operate in rings with impunity. For example, in March 2001, a U.S. federal grand jury charged three men for their participation in a ring of fraudulent bidding in hundreds of art auctions on eBay

(see Schwartz and Dobrzynski [1]). The men created more than 40 User IDs on eBay using false registration information.

Trevathan and Read [4] presented an algorithm (referred to as the shill score or SS algorithm) to detect the presence of shill bidding in online English auctions. The algorithm observes bidding patterns over a series of auctions, and gives each bidder a *shill score* to indicate the likelihood that they are engaging in shill behaviour. The algorithm is able to accurately identify those with suspicious behaviour in the case where there is only one shill bidder. However, as the example above illustrates, there are situations where multiple shills are in collusion with each other.

Collusive shill bidding refers to the instance where two or more shills work together. Collusion makes it more difficult to determine whether shilling is occurring in an auction, and which bidders are responsible for the shill bids. This is because colluding shills can distribute the work evenly among each other to collectively reduce their shill scores. For example, a group of shills can take turns at submitting bids, and/or alternate at participating in an auction. This has the effect of making an individual shill appear to be a regular bidder.

This paper proposes a method for detecting colluding shill bidders. Several *collusion ratings* are presented that indicate whether collusion is occurring, and which bidders are most likely to be in collusion with each other. The collusion ratings look for typical collusive behaviour among bidders, and form a bidder's *collusion score*. The group of shills is faced with a choice of either, a) acting as a group and being detected by the collusion score, or b) acting individually (like a single shill), and being detected by the SS algorithm. The collusion score has been tested on simulated auction data to gauge its effectiveness in detecting colluding shills.

This paper is organised as follows: Background on general shill behaviour is discussed in Section II. Section III describes shill detection methods and provides an overview of the SS algorithm. Section IV extends the SS algorithm to detect collusive shill behaviour, and identify which bidders are in collusion with each other. Section V describes how the proposed collusion score performs on simulated auction data with colluding shills. Section VI provides concluding remarks.

II. SHILL BEHAVIOUR

This section provides an insight into general shill behaviour. It describes a shill's mindset, characteristics and strategies, and presents an example of shill behaviour in an auction. There are

¹<http://www.ebay.com>

²<http://www.ubid.com>

other possible behaviours that can be considered as shilling. However, we believe that this section outlines the optimal or ‘best’ strategies for a shill to take in terms of obtaining the highest shilling profit, whilst avoiding winning an auction.

A. Shill Mindset

The main goal for shilling is to artificially inflate the price for the seller beyond what legitimate bidders would otherwise require to win the item. The pay-off for the seller is the difference between the final price and the uninflated price.

A shill’s goal is to lose each auction. A shill is not constrained by a budget, but rather a profit margin. If the shill wins, the item is resold in a subsequent auction. However, there is a limit on how many times this can be done. For each auction won, the seller incurs auction listing fees and is required to invest more time. Continual wins erode the profit from shilling on the item.

The shill faces a dilemma for each bid they submit. Increasing a bid could marginally increase the revenue for the seller. However, raising the price might also result in failure if it is not outbid before the auction terminates. The shill must decide whether to take the deal or attempt to increase the pay-off.

On the contrary, a bidder’s goal is to win. A bidder has a finite budget and is after the lowest price possible. Increasing a bid for a legitimate bidder decreases the money saved, but increases the likelihood of winning.

B. Shill Characteristics and Strategies

The following outlines typical shill characteristics assuming we are using an auction that terminates at a set expiration time:

1) A shill usually bids exclusively in auctions only held by one particular seller.

2) A shill tends to have a high bid frequency. An aggressive shill will continually outbid legitimate bids inflating the final price until the seller’s expected pay-off for shilling has been reached, or until the shill risks winning the auction (e.g., near the termination time or during slow bidding).

3) A shill has few or no winnings.

4) It is advantageous for a shill to bid within a small time period after a legitimate bid. Generally a shill wants to give legitimate bidders as much time as possible to submit a new bid before the closing time of the auction.

5) A shill usually bids the minimum amount required to outbid a legitimate bidder. If the shill bids an amount that is much higher than the current highest bid, it is likely that legitimate bidders will be deterred from bidding.

6) A shill’s goal is to try and stimulate bidding. As a result, a shill will tend to bid more near the beginning of an auction. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the end of an auction is risky as the shill could accidentally win.

We refer to the most extreme shill bidding strategy as *aggressive shilling*. An aggressive shill continually outbids everyone thereby driving up the price as much as possible. This strategy often results in the shill entering many bids.

TABLE I
AN EXAMPLE AUCTION WITH ONE SHILL

Bid #	b_{id}	Price	Time
15	b_1	\$ 33	20:03
14	b_2	\$ 32	12:44
13	b_1	\$ 31	12:42
12	b_2	\$ 26	5:05
11	b_1	\$ 25	5:02
10	b_2	\$ 21	2:47
9	b_3	\$ 20	2:45
8	b_2	\$ 15	1:07
7	b_1	\$ 14	1:05
6	b_2	\$ 9	0:47
5	b_3	\$ 8	0:45
4	b_2	\$ 6	0:20
3	b_3	\$ 5	0:19
2	b_2	\$ 2	0:06
1	b_1	\$ 1	0:05

In contrast, a shill might only introduce an initial bid into an auction where there has been no prior bids with the intent to stimulate bidding. This behaviour is a common practice in traditional and online auctions. However, most people typically do not consider it fraudulent. Nevertheless it is still shilling, as it is an attempt to influence the price by introducing spurious bids. We refer to this as *benign shilling* in the sense that the shill does not continue to further inflate the price throughout the remainder of the auction. A benign shill will typically make a “one-off” bid at or near the beginning of the auction.

Regardless of the strategy employed, a shill will still be a bidder that often trades with a specific seller but has not won any auctions. A shill’s strategy is also affected by the value of the current bid in relation to the reserve price. For example, once bidding has reached the reserve price it becomes more risky to continue shilling. This is conditional on whether the reserve is a realistic valuation of the item that all bidders share.

C. Shill Example

Table I illustrates an example auction with three bidders. Each bidder is denoted as b_1 , b_2 and b_3 respectively. Bidders b_1 and b_3 are legitimate, whereas b_2 is a shill. b_2 engages in aggressive shill behaviour by outbidding a legitimate bid by the minimal amount required to stay ahead and within a small time period of the last bid. b_2 ’s bids force the other bidders to enter higher bids in order to win. If b_2 was not participating in this auction, b_1 would have only needed to pay \$21 in order to win. Instead b_2 caused b_1 to pay \$33, thus the shill has inflated the price by \$12.

In this example, b_2 exhibits the typical shill behaviour described above. This is evidenced by: 1. High frequency of bids (i.e., b_2 has submitted more bids than both the other bidders); 2. Has not won the auction despite the high number of bids; 3. Quick to bid after a legitimate bidder; and 4. Only bids the minimal amount to stay in front.

III. SHILL DETECTION

Until recently, there was limited coverage on exactly how to detect shill bidding. Basic attempts have been proposed by Shah *et al* [2] and Rubin *et al* [3]. Trevathan and Read [4] propose a solution (i.e., the SS algorithm), that observes

bidding patterns over a series of auctions for a particular seller, looking for the shilling behaviour outlined in Section II. The goal is to obtain statistics regarding a bidder’s conduct, and deduce a measure called a *skill score*, that indicates the likelihood that s/he is engaging in shill behaviour.

The SS algorithm targets core shilling strategies. A shill that deviates too far from these characteristics is less effective, and won’t significantly alter the auction outcome. This approach acts as both a detection mechanism and a deterrent to shill bidders. To avoid detection, a shill must behave like a normal bidder, which essentially stops them shilling.

The SS algorithm basically works as follows (see Trevathan and Read [4] for further details): A bidder i , is examined over m auctions held by the same seller for the behaviour outlined in Section II. Each characteristic of shill behaviour is given a rating, which is combined to form the bidder’s skill score. The skill score gives a bidder a value between 0 and 10. The closer the skill score is to 10, the more likely that the bidder is a shill. The algorithm’s goal is to determine which bidder is most inclined to be the shill out of a group of n bidders. The skill behavioural ratings are calculated as follows:

- α Rating - Percentage of auctions bidder i has participated in.
- β Rating - Percentage of bids bidder i has made out of all the auctions participated in.
- γ Rating - Normalised function based on the auctions bidder i has won out of the auctions participated in.
- δ Rating - Normalised inter bid time for bidder i out of the auctions participated in.
- ϵ Rating - Normalised inter bid increment for bidder i out of the auctions participated in.
- ζ Rating - Normalised time bidder i commences bidding in an auction.

Each rating is between 0 and 1, where the higher the value, the more suspicious the bidder. A bidder’s skill score (denoted as SS) is calculated as the weighted average of these ratings:

$$SS = \frac{\theta_1\alpha + \theta_2\beta + \theta_3\gamma + \theta_4\delta + \theta_5\epsilon + \theta_6\zeta}{\theta_1 + \theta_2 + \theta_3 + \theta_4 + \theta_5 + \theta_6} \times 10$$

where θ_i , $1 \leq i \leq 6$, is the weight associated with each rating. [4] provides details and justifications for selection of weight values. If a bidder wins an auction, then his/her α , β , δ , ϵ and ζ ratings are 0 for the particular auction.

IV. COLLUDING SHILLS

The SS algorithm considers only the basic scenario with **one shill**. In the case outlined in Section I, there were three sellers which used 40 different aliases (40 shills in effect). The sellers understood that there was less chance that they would get caught if they utilised multiple names to take alternating turns at shilling. This makes it more difficult for authorities to determine which bidders are shills, as collusive behaviour allows shills to appear as more regular bidders.

In terms of the SS algorithm, introducing multiple shills makes the task of detection much harder. This is because colluding shills can engage in more sophisticated strategies

in an attempt to thwart the detection algorithm. This section discusses these strategies and presents methods by which the SS algorithm can be extended to detect colluding shills. We restrict our attention to the case where there is only **one** seller which controls multiple shills. That is, we are not investigating strategies involving multiple sellers.

In some cases, geographical proximity can be an indication of collusion if there are several shills within a close area that participate in the auction. For example, in the shill case, two of the men were from California and the other was from Colorado. However, this is not a reliable indicator of shilling and may raise privacy concerns, as it requires examining the registration database for such relationships. Checking suspect colluders’ geographical proximity would generally only occur once the SS-collusion algorithm has been run, thus ensuring there is strong evidence for doing so.

As stated previously, the main goal of shilling is to *drive up the price of an item*. In the situation where there is only one shill, the shill’s secondary goal is to attempt to do this in such a manner that it *minimises his/her skill score*. When there is more than one shill, there are particular strategies that the group (of shills) can engage in to influence some factors contributing to their individual skill scores. Therefore the group’s collective goal (secondary to shilling) is to *minimise its member’s skill scores*.

Despite being able to use more complicated strategies, the group as a whole must still conform to certain behaviour in order to be effective as a shill. With regard to the SS algorithm, all that shills can do by colluding is to reduce their α , β and ζ ratings. The γ , δ and ϵ ratings are still indicative of shill bidding. For example, none of the colluding shills will ever win an auction. Furthermore, it is still in the group’s interests to bid quickly, and by minimal amounts to influence the selling price. Therefore, inter bid times and increments will be consistent for shills.

There appear to be three possible strategies that can be employed by colluding shills. The following example illustrates the first strategy. Here there are two colluding shills (s_1 and s_2). Each shill takes alternating turns at bidding, i.e., s_1 bids, then s_2 bids, then s_1 bids again, etc. This behaviour is shown

TABLE II

AN EXAMPLE AUCTION WITH TWO SHILLS ALTERNATING BIDS

Bid #	b_{id}	Price	Time
1	b_1	1	0:05
2	b_2	2	0:06
3	b_1	5	0:19
4	b_3	6	0:20
5	b_1	8	0:45
6	b_2	9	0:47
7	b_1	14	1:05
8	b_3	15	1:07
9	b_1	20	2:45
10	b_2	21	2:47
11	b_1	25	5:02
12	b_3	26	5:05
13	b_1	31	12:42
14	b_2	32	12:44
15	b_1	35	20:03

in Table II. Here there are three bidders denoted b_1, b_2 and b_3 respectively. b_1 is a legitimate bidder, but b_2 and b_3 are skills. b_2 and b_3 take alternating turns at outbidding b_1 .

This strategy has the effect of lowering b_2 and b_3 's β ratings (i.e., the number of individual skill bids in an auction). We refer to this strategy as the *alternating bid* strategy.

The second strategy is for colluding skills to take turns at shilling for a particular auction. For example, given two auctions, one skill will bid exclusively in the first auction, while the other skill bids only in the second auction. This strategy lowers the skills' α ratings (i.e., number of auctions participated in), but does not affect their β ratings. We refer to this strategy as the *alternating auction* strategy.

The third strategy is to use a combination of the alternating bid and alternating auction strategies. This can be used to alter the group's α and β ratings between the two extremes. We refer to this strategy as the *hybrid* strategy. An example of a hybrid strategy would be for skills s_1 and s_2 to alternately bid in auction₁, s_3 and s_4 alternately bid in auction₂, then s_1 and s_3 alternately bid in auction₃, etc. This continues until all combinations of bidders have been used, and then the process repeats. In reality, colluding skills would probably employ a hybrid strategy.

The following notation is used throughout this paper:

Let $L = \{1, \dots, \ell\}$ be the set of bidder numbers; $|L| = \ell$

Let $M = \{1, \dots, m\}$ be the set of auction numbers; $|M| = m$

Let $\mathcal{B} = \{b_1, \dots, b_\ell\}$ be the set of bidders and $\mathcal{S} = \{s_1, \dots, s_{\ell'}\}$, $\mathcal{S} \subset \mathcal{B}$, be the set of skills, where $\ell' < \ell$.

A. Alternating Bid Strategy

If the alternating bid strategy is employed, it is in the group's best interests to evenly alternate. If a particular skill bids more than other skills, this will increase that individual's β rating and hence their skill score. This violates the group's collective goal of minimising all its member's scores. There are rules that govern the number of bids a group of skills can submit during an auction. We will examine some of the possible scenarios.

Table III illustrates a case where there are two bidders and one of them is a skill. b_1 must bid $\lceil n_j/2 \rceil$ times to win. The skill can submit at most $\lfloor n_j/2 \rfloor$ bids, otherwise they will win the auction (and fail as a skill).

Table IV illustrates a case where there are three bidders and one of them is a skill. Consider the worst case scenario for the skill in terms of the SS algorithm where they must bid after

TABLE III

AUCTION WITH TWO BIDDERS - ONE BIDDER (b_1), ONE SKILL (s_1)

n_j	b_1	s_1	bid sequence	b_1	s_1
3	2	1	b_1, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$
4	2	2	s_1, b_1, s_1, b_1	$n_j/2$	$n_j/2$
5	3	2	b_1, s_1, b_1, s_1, b_1	$\lceil n_j/2 \rceil$	$\lfloor n_j/2 \rfloor$
6	3	3	$s_1, b_1, s_1, b_1, s_1, b_1$	$n_j/2$	$n_j/2$

TABLE IV

AUCTION WITH THREE BIDDERS - TWO BIDDERS (b_1, b_2), ONE SKILL (s_1)

n_j	b_1	b_2	s_1	bid sequence	b_1	b_2	s_1
3	1	1	1	b_1, s_1, b_2	$\lceil n_j/2 \rceil$		$\lfloor n_j/2 \rfloor$
4	1	1	2	s_1, b_1, s_1, b_2	$n_j/2$		$n_j/2$
5	2	1	2	b_1, s_1, b_2, s_1, b_1	$\lceil n_j/2 \rceil$		$\lfloor n_j/2 \rfloor$
6	2	1	3	$s_1, b_1, s_1, b_2, s_1, b_1$	$n_j/2$		$n_j/2$

TABLE V

AUCTION WITH THREE BIDDERS - ONE BIDDER (b_1), TWO SKILLS (s_1, s_2)

n_j	b_1	s_1	s_2	bid sequence	b_1	s_1	s_2
4	2	1	1	s_1, b_1, s_2, b_1	$n_j/2$		$n_j/2$
5	3	1	1	b_1, s_1, b_1, s_2, b_1	$\lceil n_j/2 \rceil$		$\lfloor n_j/2 \rfloor$
6	3	2	1	$s_1, b_1, s_2, b_1, s_1, b_1$	$n_j/2$		$n_j/2$
7	4	2	1	$b_1, s_1, b_1, s_2, b_1, s_1, b_1$	$\lceil n_j/2 \rceil$		$\lfloor n_j/2 \rfloor$
8	4	2	2	$b_1, s_1, b_1, s_2, b_1, s_1, b_1, s_2$	$n_j/2$		$n_j/2$

every legitimate bid. This situation results in the maximum β rating a bidder could possibly get. As in the previous example, the skill can submit at most $\lfloor n_j/2 \rfloor$ bids. In reality, a skill will not normally need to bid this many times, as there are would be situations where two or more legitimate bidders might outbid another bid before a skill has time to respond. However, we must remember that is the worst case scenario (i.e., the skill must bid after every legitimate bid). In contrast, the best case scenario would be where the skill wouldn't have to bid at all, because the legitimate bidders had pushed the price beyond the seller's expected pay-off for shilling.

Given the worst case scenario with n legitimate bidders and one skill, regardless of which bidder bids, collectively the legitimate bidders must bid $\lceil n_j/2 \rceil$ times to win. This is shown in Table IV, where bidders b_1 and b_2 combined must outbid the skill, s_1 .

In terms of colluding skills, increasing the number of skills decreases the number of times each skill must bid (also decreasing an individual's β rating). This alleviates the worst case scenario for shilling. Given ℓ' skills (where $\ell' < \ell$), the group can split the blame by dividing the number of bids submitted amongst themselves. An individual skill can submit as little as $\lfloor n_j/2 \rfloor / \ell'$ bids in an auction. This is shown in Table V with one bidder and two skills.

In the case of *maximum collusion* (i.e., all bidders are skills except for one, $\ell' = \ell - 1$), to win the legitimate bidder is forced into submitting the most bids for an individual in the auction (i.e., $\lceil n_j/2 \rceil$ bids). If the legitimate bidder wins, then his/her β rating will be 0 for the auction. However if the bidder does not win, his/her β rating will be significantly higher than all the other bidders (i.e., the skills), as a result of having to compete with the skill bids. The later scenario becomes common as more legitimate bidders are added. In this situation a bidder's γ rating (number of wins), will be a significant factor in determining whether they really are a skill. This example shows how colluding skills can lower their individual skill scores while shifting attention to innocent bidders.

The major question maximum collusion poses is, "how do we detect and deter collusion among skill bidders to




Fig. 1. Example Collusion Graph

ensure that the SS algorithm does not incriminate a legitimate bidder?" We propose a solution to the problem, which we refer to as the *collusion graph*.

- Collusion Graph -

The collusion graph indicates which bidders are likely to be in collusion with each other. There are two different forms of the collusion graph based on whether skills use the alternating bid, or the alternating auction strategy. Two new collusion ratings are introduced, η and θ , that measures a bidder's conduct in terms of collusive behaviour. These ratings **are not** used in the calculation of the skill score. Instead, the ratings supplement the skill score, and serve to bring possible colluders to the attention of the SS algorithm. We will discuss each rating in turn.

To detect colluding groups employing the alternating bid strategy, bidders are represented as a graph $G = (V, E)$. V is the set of bidders, and E is an edge between two bidders, indicating that they have both participated in the same auction. The goal of the collusion graph is to find a subset, $C \subset V$, which contains the bidders that are most likely to be in collusion with each other.

A bidder initially has no edges connecting it to other bidders (i.e., the set E is empty). Given two bidders, $v_i, v_j \in V$, $i \neq j$, that participate in the same auction, an edge $e_{i,j}$ is added to E that connects these two bidders together. If bidders v_i, v_j participate in more than one auction simultaneously, weights are added to the edge $e_{i,j}$ to indicate the number of auctions they were present in together.

The idea is that G will form a tree connecting colluding bidders together. The higher the edge weighting between two bidders, the greater the likelihood that the two bidders are in collusion with each other. Figure 1 gives an example of a collusion graph. Here there are two colluding bidders. In general, skills will have the most number of edges (i.e., highest degree), and higher edge weightings than legitimate bidders.

Given a node with degree k , each edge weight is denoted as w_j , $1 \leq j \leq \ell$. The base collusion rating, η'_i , for a bidder i is calculated as the sum of the edge weights for the degree of the node:

$$\eta'_i = \sum_j^k w_j$$

To calculate bidder i 's normalised collusion rating η_i , we first need to find the maximum and minimum base collusion ratings for the graph. These are denoted as η^{max} and η^{min}




Fig. 2. Potential Colluding Bidders




Fig. 3. Example Collusion Graph

respectively. The normalised collusion rating is calculated as:

$$\eta_i = \frac{\eta'_i - \eta^{min}}{\eta^{max} - \eta^{min}}$$

where $0 \leq \eta_i \leq 1$.

Colluding skills will have similar η collusion ratings. This value is initially used to bring suspected colluders to the attention of the SS algorithm, (but is not used in the calculation of the skill score). Given a bidder with a high collusion rating, we can then observe his/her γ , δ and ϵ ratings, which will provide a strong indication whether the bidder is engaging in collusive behaviour. Consistent features for the group of skills will be their γ , δ and ϵ ratings.

Figure 2 shows the previous example of the collusion graph after the algorithm has been run. The bold nodes indicate the suspect bidders. The η ratings for these bidders is 1. The next highest collusion rating for a bidder is 0.27, which is significantly lower than the three suspected colluders. This example shows that the collusion graph is able to immediately single out suspect bidders. These bidders can then be further investigated based on their γ , δ and ϵ ratings.

Figure 3 illustrates an example collusion graph with two skills and three bidders. In this example there are three auctions. Each bidder participates in exactly one auction, whereas the skills participate in all three auctions. The bidding sequences for each auction are $(s_1, b_1, s_2, b_1, s_1, b_1)$, (s_2, b_2, s_1, b_2) and $(s_2, b_3, s_1, b_3, s_2, b_3)$ respectively. The right side of Figure 3 shows that after the algorithm has been run, the colluding skills strongly emerge as the most suspicious.

The collusion graph can be represented as an adjacency matrix. Each entry in the matrix lists the edge weighting for a pair of bidders. The collusion graph from Figure 3 gives the following adjacency matrix:

	s_1	s_2	b_1	b_2	b_3
s_1	—	3	1	1	1
s_2	3	—	1	1	1
b_1	1	1	—	0	0
b_2	1	1	0	—	0
b_3	1	1	0	0	—

Base collusion ratings for each bidder can be obtained by summing the entries in a bidder's column.

Colluding shills using the alternating bid strategy will have almost exactly the same α rating. The collusion graph reflects this in the edge weights between bidders. Colluding shills will also have participated with a similar number of legitimate bidders. The collusion graph reflects this by the degree of a bidder's node. Colluding shills will have almost exactly the same η rating.

However, in this form the η rating does not significantly bind groups of bidders together. There are a lot of innocent bidders that are incorporated into the collusion graph by chance. These bidders will tend to have the lower η ratings. Less significant collusion structures can be pruned from the collusion graph, by using the observation that colluding shills using the alternating bid strategy will have almost exactly the same β rating.

Given two suspect colluding bidders, we can further compare them based on their β ratings. We refer to this as a *binding factor*. The binding factor, $\phi_{i,j}^\beta$, gives bidders i and j a value between 0 and 1, based on how similar their β ratings are. A binding value of 1 indicates a strong binding (high likelihood of collusion), whereas a binding value of 0 indicates a weak binding.

Let β_i and β_j denote the β ratings for bidders i and j respectively. If $\beta_i = \beta_j$, then the binding factor, $\phi_{i,j}^\beta$, is 1. If this is not the case, then the binding factor is given a decreasing value based on the percentage difference between the two values. The binding factor for two suspect bidders is calculated as:

$$\phi_{i,j}^\beta = \begin{cases} 1 & \text{if } \beta_i = \beta_j \\ \beta_i/\beta_j & \text{else if } \beta_i < \beta_j \\ \beta_j/\beta_i & \text{otherwise} \end{cases}$$

where $0 \leq \phi_{i,j}^\beta \leq 1$.

Consider the maximum collusion scenario held over m auctions. The legitimate bidder will have high α and β ratings (i.e., high participation rate and number of bids). The shills will have the same α ratings as the bidder, but low β ratings. The collusion graph will initially incriminate the legitimate bidder with the shill's group (due to the high participation rate). However, the legitimate bidder will have a high β rating, as they must submit more bids than the group of shills to win (i.e., $\lceil n_j/2 \rceil$ bids).

In this case, the binding factor, $\phi_{i,j}^\beta$ can be used to compare the bidder to any of the shills. $\phi_{i,j}^\beta$ will be low for the bidder, (i.e., s/he does not bind to any of the shills). Whereas $\phi_{i,j}^\beta$ will be high for the colluding group. Therefore $\phi_{i,j}^\beta$ serves to exonerate the legitimate bidder. Furthermore, if the bidder has actually managed to win any auctions, his/her γ rating will be low and thus provide a strong indication of integrity.

In order to determine which bidders are the most likely to be in collusion, a *collusion score* (denoted as CS^η) is used. CS^η combines the η collusion rating, binding factor and other aspects of the shill score ratings, to form an index which measures the extent of collusive behaviour a bidder exhibits in terms of the alternating bid strategy.

The set of bidders \mathcal{B} , is partitioned into disjoint sets of suspected colluding bidders. This is done based on how

similar bidders' η ratings are to other bidders. Let $\mathcal{C} = \{C_1, C_2, \dots, C_k\}$ denote the set of all sets of suspicious bidders. Initially, the first bidder, i is assigned to C_1 . Bidder j ($i \neq j$) is added to C_1 if $\eta_i = \eta_j \pm \lambda$, where λ is an error factor. This process is repeated for all remaining bidders until every bidder is assigned to a set in \mathcal{C} . A bidder that is not suspicious will be assigned to a singleton set (i.e., s/he is the only bidder in the set). Note that this is probably not the 'best' method for grouping bidders, as some bidders whose η ratings fall within the range of $\eta_i \pm \lambda$ might miss out on being grouped with more similar bidders later in the process (i.e., those with a closer η rating). A more optimal grouping method remains the focus of future work.

Once two (or more) bidders have been singled out based on the similarity of their collusion ratings, CS^η checks the bidder's binding ratings. ϕ_σ^β is the average binding factor for a bidder with the other suspect colluding group members. For a singleton set $\phi_\sigma^\beta = 0$. A potential shill will have a low win rate, high bid rate, and low bid increment. If $\eta_i > 0.5$, CS^η is calculated as the average of the bidder's γ , δ , ϵ and η ratings, and the average binding factor ϕ_σ^β . A bidder with a shill score of zero is excluded. Bidder i 's CS^η is calculated as:

$$CS_i^\eta = \frac{\gamma + \delta + \epsilon + \eta + \phi_\sigma^\beta}{5} \times 10$$

where $0 \leq CS^\eta \leq 10$.

The ζ rating is not included in calculating CS^η . This is because for an individual auction, the shill that bids first will have a higher ζ rating than the later shills. As a result, over m auctions, each shill must also alternate at bidding first. This ensures that over time each of the group members will have a similar ζ rating. A binding factor could be used for two suspect bidders based on their ζ ratings. However, this is a less reliable collusion indicator.

B. Alternating Auction Strategy

The alternating auction strategy is optimal for a colluding group, if the number of shills, ℓ' , is equal to the number of auctions held, m . In the shill case from Section I, there were 1,100 auctions, however, the shills only had 40 aliases. In this case, given m auctions, the group can evenly reduce its member's α ratings if each shill participates in m/ℓ' auctions. However, there is a certain amount of effort required to create new aliases and/or recruit new shills. This reduces this strategy's effectiveness, especially when m is large.

- Dual Collusion Graph -

The alternating auction strategy requires a different approach. The idea is to look for bidders that have never competed against each other in an auction. In this situation, the "dual" of the collusion graph is used. In the dual collusion graph, the nodes that have edges in the regular graph are removed, and edges are added between nodes that did not have edges in the regular graph. Edge weights from the regular graph have no meaning, as either two bidders have participated in at least one auction together, or they haven't. All edge

TABLE VI
SHILL SCORE AND COLLUSION BEHAVIOUR PROFILES

Rating		Single Skill	Alternating Bid	Alternating Auction	Hybrid
α	participation rate	high	high	low	low-medium
β	bid rate	high	low	high	low-medium
γ	loss rate	high	high	high	high
δ	bid frequency	high	high	high	high
ϵ	size of bid	high	high	high	high
ζ	early participation	high	medium	high	medium-high
η	alternating bid		high		medium-high
θ	alternating auction			high	low-medium
ϕ_{σ}^{β}	binding factor		high		medium-high
ϕ_{σ}^{α}	binding factor			high	medium-high




Fig. 4. Example Dual Collision Graph

weights from the regular graph are discarded. All edges in the dual graph have a weight of 1.

Figure 4 shows an example of the dual collision graph involving two bidders and two skills. There are two auctions. The skills participate in one auction each, whereas the regular bidders participate in both auctions. The bidding sequences for each auction are (s_1, b_1, s_1, b_2) and (s_2, b_1, s_2, b_2) respectively. The adjacency matrix for this collision graph is:

$$\begin{array}{ccccc}
 & s_1 & s_2 & b_1 & b_2 \\
 s_1 & - & 0 & 1 & 1 \\
 s_2 & 0 & - & 1 & 1 \\
 b_1 & 1 & 1 & - & 2 \\
 b_2 & 1 & 1 & 2 & -
 \end{array}$$

The dual collision graph shows an association between s_1 and s_2 . The adjacency matrix for the dual collision graph is:

$$\begin{array}{ccccc}
 & s_1 & s_2 & b_1 & b_2 \\
 s_1 & - & 1 & 0 & 0 \\
 s_2 & 1 & - & 0 & 0 \\
 b_1 & 0 & 0 & - & 0 \\
 b_2 & 0 & 0 & 0 & -
 \end{array}$$

θ is a collusion rating that indicates whether a bidder may be engaging in the alternating auction strategy. This is obtained from the dual collision graph in a similar manner to the η rating. Given a node with degree k , each edge weight is denoted as w_j , $1 \leq j \leq \ell$. The base collusion rating, θ'_i , for a bidder i is calculated as the sum of the edge weights for the node's degree:

$$\theta'_i = \sum_j^k w_j$$

To calculate a bidder's normalised collusion θ rating we first need to find the maximum and minimum base collusion ratings for the graph. These are denoted as θ^{max} and θ^{min}

respectively. The normalised collusion rating is calculated as:

$$\theta_i = \frac{\theta'_i - \theta^{min}}{\theta^{max} - \theta^{min}}$$

where $0 \leq \theta_i \leq 1$.

Given two suspect colluding bidders, we can further compare them based on their α ratings. The binding factor, $\phi_{i,j}^{\alpha}$, gives bidders i and j a value between 0 and 1, based on how similar their α ratings are. A binding value of 1 indicates a strong binding (high likelihood of collusion), whereas a binding value of 0 indicates a weak binding.

Let α_i and α_j denote the α ratings for bidders i and j respectively. If $\alpha_i = \alpha_j$, then the binding factor, $\phi_{i,j}^{\alpha}$, is 1. If this is not the case, then the binding factor is given a decreasing value based on the percentage difference between the two values. The binding factor for two suspect bidders is calculated as:

$$\phi_{i,j}^{\alpha} = \begin{cases} 1 & \text{if } \alpha_i = \alpha_j \\ \alpha_i/\alpha_j & \text{else if } \alpha_i < \alpha_j \\ \alpha_j/\alpha_i & \text{otherwise} \end{cases}$$

where $0 \leq \phi_{i,j}^{\alpha} \leq 1$.

In order to determine which bidders are the most likely to be in collusion, a *collusion score* (denoted as CS^{θ}) is used. CS^{θ} combines the θ collusion rating, binding factor and other aspects of the skill score ratings, to form an index which measures the extent of collusive behaviour a bidder exhibits in terms of the alternating auction strategy.

The set of bidders \mathcal{B} , is partitioned into disjoint sets of suspected colluding bidders. This is done in a similar manner as the alternating bid strategy. Once two (or more) bidders have been singled out based on the similarity of their collusion ratings, CS^{θ} checks the bidder's binding ratings. ϕ_{σ}^{α} is the average binding factor for a bidder with the other suspect colluding group members. For a singleton set $\phi_{\sigma}^{\alpha} = 0$. A potential skill will have a high bid rate, low bid increment, and generally commences bidding early in an auction. If $\theta_i > 0.5$, CS^{θ} is calculated as the average of the bidder's δ , ϵ , ζ , and θ ratings, and the average binding factor ϕ_{σ}^{α} . Again, a bidder with a skill score of zero is excluded. Bidder i 's CS^{θ} is calculated as:

$$CS_i^{\theta} = \frac{\delta + \epsilon + \zeta + \theta + \phi_{\sigma}^{\alpha}}{5} \times 10$$

where $0 \leq CS^{\theta} \leq 10$.

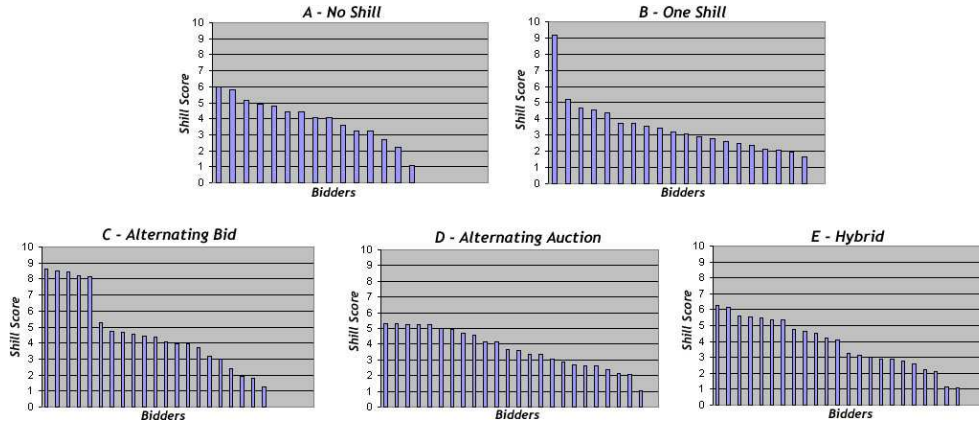


Fig. 5. Shill Scores for Simulated Auctions with Differing Collusion Strategies

Note that the γ rating is excluded as it is a function of the number of auctions participated in, versus the number of winnings. This is unreliable if the colluding group has engaged in an alternating auction strategy, as the α rating will be low, thus affecting the γ rating. The alternating auction strategy has no ramifications for the ζ rating (i.e., the time a bidder commences bidding). That is, for each auction the shill will behave in a typical shill-like manner (i.e., bidding early). This means that the ζ rating is still indicative of a shill.

C. Hybrid Strategy

A hybrid strategy combines both the alternating bid and alternating auction strategies. By doing so, the colluding group can lower each member's α and β ratings. The optimal hybrid strategy requires the coalition to choose $\binom{\ell'}{p} = m$, where ℓ' is the total number of shills, p is the number of shills to be used each auction, and m is the total number of auctions.

Table VI shows the shill score ratings for each shilling strategy. Of particular interest here is the comparison between a single shill's profile and a group of colluding shills using a hybrid strategy. At most the group can only minimally influence their shill scores. However, if they do so they will be detected by the collusion graph. The only way to avoid detection is for one of the group members to have a shill score higher than the others (thus conforming to the profile of a single shill).

A hybrid strategy can be detected using a combination of the η and θ ratings. The ideal situation for the group is to keep the α and β ratings even for all colluders. However, the binding factors $\phi_{i,j}^\beta$ and $\phi_{i,j}^\alpha$, prevent the group from evenly distributing their work. As doing so will result in the colluding group having high binding factors. This forces some shills to bid more, and/or participate in more auctions than other shills (thus raising their individual α and β ratings). An extreme hybrid strategy would entail using each shill once only. However, there are practical limitations on doing this.

To detect shills employing a hybrid strategy, we introduce a third collusion score (denoted as CS^h). CS^h examines the η ratings, and the binding factors for bidders that fit the hybrid profile shown in Table VI. θ is not used, as this value is typically low for the hybrid strategy. The δ and ϵ ratings will

still be indicative of a shill bidder. However the α , β and γ ratings can not be relied on. At best the hybrid strategy only achieves a moderate reduction in the group's ζ ratings. Again, a bidder with a shill score of zero is excluded. Bidder i 's CS^h is calculated as:

$$CS_i^h = \frac{\delta + \epsilon + \eta + \phi_\sigma^\beta + \phi_\sigma^\alpha}{5} \times 10$$

where $0 \leq CS^h \leq 10$.

V. PERFORMANCE


This section describes how the collusion algorithm performs on simulated auction data generated using autonomous bidding agents. The Research Auction Server (RAS) at James Cook University is an online auction server used for conducting research into privacy and security issues in online auctions (see [6]). The SS algorithm and collusion graph have been implemented on and tested using RAS. Elements of the experimental setup include:

Zero Intelligence (ZI) Agent - An agent designed to simulate an ordinary bidder in an auction. It is assigned a random amount, which it tries to submit as a proxy bid at a random time throughout the auction. The agent's price is generated randomly according to a uniform distribution.

Simple Shill Bidding Agent - An agent designed to insert fake bids into the auction in order to inflate the price for the seller. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding (e.g., during slow bidding, or near the auction's end). See [7] for implementation specific details regarding the simple shill agent.

Auction - A software simulated auction. The auction has a start and end time. The bidding agents can submit bids during this time, where the auction outcome depends on the agents' actions. All auctions are English auctions with proxy bidding.

Four main test types were conducted. Each test consisted of ten auctions with 20 ZI bidders. All ZI bidders were given the opportunity to participate in any auction. However not all ZI agents can participate in every auction. This is because a ZI agent is unable to participate in a particular auction if the current bid exceeds its proxy bid at the stage in the



Userid	α	β	γ	δ	ϵ	ζ	Shill Score	η	Group	ϕ_σ^α	Collusion Score
SHILL_a	0.7	0.09	0.86	0.76	0.97	0.73	7.13	1	10	0.99	8.54
SHILL_b	0.7	0.09	0.86	0.75	0.96	0.6	7	0.99	10	0.99	8.51
SHILL_f	0.6	0.09	0.83	0.88	0.97	0.75	6.79	0.9	12	0.91	8.38
SHILL_g	0.6	0.1	0.83	0.79	0.96	0.46	6.44	0.83	11	0.95	8.26
SHILL_e	0.6	0.1	0.83	0.78	0.96	0.55	6.52	0.83	11	0.95	8.25
SHILL_d	0.6	0.09	0.83	0.78	0.96	0.6	6.74	0.85	11	0.95	8.24
SHILL_g	0.6	0.09	0.83	0.74	0.99	0.72	6.66	0.8	11	0.96	8.24
SHILL_h	0.6	0.09	0.83	0.77	0.95	0.73	6.66	0.85	11	0.95	8.22
SHILL_i	0.6	0.08	0.83	0.76	0.96	0.69	6.61	0.92	12	0.95	8.21
SHILL_j	0.6	0.08	0.83	0.68	0.95	0.6	6.45	0.92	12	0.95	8.05
agent'o	0	0	0	0	0	0	0	0	0	0	0
agent'r	0.1	0.24	0	0.98	0.97	0.51	2.86	0.48	3	0.69	0
agent't	0.1	0.07	0	0.68	0.9	0.24	2.15	0.28	8	0.71	0
agent'e	0.2	0.07	0.5	1	0.32	0.51	3.65	0.27	8	0.72	0
agent's	0.1	0.05	0	0.9	0.92	0.14	2.23	0.45	5	0.1	0
agent'o	0.1	0.09	0	1	0.44	0.75	2.40	0.1	9	0	0
agent'f	0.3	0.29	0.67	0.92	0.92	0.65	5.27	0.48	3	0.65	0
agent'g	0.1	0.07	0	1	0	0.78	2.09	0.17	6	0.15	0
agent'e	0.1	0.5	0	0.83	0.96	0.55	2.99	0.42	5	0.1	0
agent'd	0	0	0	0	0	0	0	0	0	0	0
agent'b	0.4	0.24	0.75	0.71	0.48	0.37	4.97	0.59	2	0	0
agent'c	0.3	0.13	0.67	0.66	0.34	0.59	4.31	0.49	3	0.51	0
agent'h	0	0	0	0	0	0	0	0	0	1	0
agent'i	0.2	0.44	0.5	0.73	0.79	0.93	4.06	0.2	6	0.15	0
agent'm	0.2	0.09	0	1	0.66	0.43	2.8	0.34	1	0.41	0
agent'n	0.1	0.22	0	0.95	0.69	0.67	2.71	0.23	4	0	0
agent'l	0.2	0.07	0.5	0	0.8	0	2.75	0.25	8	0.71	0
agent'k	0	0	0	0	0	0	0	0	0	1	0
agent'j	0.2	0.26	0.5	0.5	0.41	0.96	3.9	0.27	8	0.26	0
agent's	0.3	0.23	0.67	0.99	0.48	0.37	4.62	0.35	1	0.41	0

Fig. 6. Example collusion scores for the alternating bid strategy

auction when it is first permitted to bid. The number of skills used varied with the type of test being conducted. The shill agent's target price was set to \$5.00 and programmed to bid in a moderately risk-adverse manner (see [7]). All tests were repeated, and standard results are given.

Claim 1 Colluding skills were able to reduce their collective shill score in relation to a single shill. As expected, the colluding skills have lower α and β ratings than a single shill, and as a result lower shill scores. Figure 5 A shows a series of auctions without shilling. The single shill's shill score was 9 (see Figure 5 B), whereas five colluding skills using the alternating bid strategy managed to reduce their shill scores to 8 (see Figure 5 C). Five colluding skills employing the alternating auction strategy achieved shill scores around 5 (see Figure 5 D). This clearly shows that the alternating auction strategy is more effective against the SS algorithm. However, as previously stated, the alternating auction strategy is more expensive in terms of the number of skills required, when compared to other strategies. Five colluding skills employing a hybrid strategy achieved a shill scores in the range of 5 to 6 (see Figure 5 E). In general, the larger the colluding group, the lower the group's shill scores.

Claim 2 The collusion graph accurately identified bidders that engaged in the alternating bid strategy. This test involved ten colluding skills employing the alternating bid strategy. Figure 6 graphically shows each bidder's collusion score. The ten skills consequently have the highest ratings. Figure 6's right side shows the shill statistics for the top ten rating bidders in terms of collusion score. The skills were divided into three distinct groups, which none of the regular bidders were members. The suspicious group members almost perfectly bind with ϕ_σ^β ranging between 0.91 and 0.99. It is interesting to note that all the legitimate bidders scored 0. Results were consistent when this test was repeated.

Claim 3 The dual collusion graph highlighted bidders that exhibited behaviour indicative of the alternating auction

strategy. This test used ten colluding skills, where each skill participated exclusively in one of the ten auctions. Figure 7 graphically shows each bidder's collusion score. In this example, 14 bidders attain a collusion score greater than 0. Out of these bidders, 11 rate higher than 7. The skills had the top ten highest consecutive ratings. Figure 7's right side shows the shill statistics for the top ten bidders in terms of the collusion score. The skills were allocated to five groups. However, several legitimate bidders were also caught up in some of these groups. This illustrates that detecting skills is not an exact science. Innocent bidders can be falsely associated with a colluding group. Likewise a shill might miss out on being assigned to a suspect group, hence avoiding detection. Suspicious group members bind strongly with ϕ_σ^α ranging between 0.80 and 1.

Claim 4 The collusion score was able to successfully single out suspect bidders employing a hybrid strategy. To validate this claim, various tests were conducted involving numerous hybrid strategies. For hybrid strategies that favoured either extreme (i.e., alternating bid or alternating auction) rated highly in the respective collusion scores. Figure 8 shows the shill and collusion scores for a hybrid strategy using $\binom{5}{2}$ combinations of skills. The five skills were the highest rating in terms of CS^h . However, similar to CS^θ , some innocent bidders are incorrectly grouped with the skills.

VI. CONCLUSIONS

This paper proposes a method for detecting colluding skill bidders. We have identified three main strategies a colluding group of skills can engage in to collectively reduce their individual shill scores. Colluding skills can be detected by examining a group of bidders for typical collusive behaviour in terms of the identified strategies. We introduce the collusion graph, which indicates relationships between groups of bidders that exhibit collusive behaviour in the form of the alternating bid strategy. The dual collusion graph is used to indicate relationships between groups of bidders that exhibit collusive behaviour in the form of the alternating auction strategy.




Fig. 7. Example collusion scores for the alternating auction strategy

Userid	α	β	γ	δ	ϵ	ζ	Shill Score	η	Group η	ϕ^s_α	Group θ	ϕ^s_θ	Collusion Score
SHILL_c	0.4	0.23	0.75	0.94	0.96	0.9	6.09	0.72	9	0.88	7	0.83	8.68
SHILL_b	0.4	0.25	0.75	0.95	0.98	0.94	6.17	0.76	9	0.83	8	0.8	8.65
SHILL_d	0.4	0.19	0.75	0.86	0.97	0.86	5.95	0.76	9	0.83	7	0.83	8.49
SHILL_a	0.4	0.21	0.75	0.95	0.97	0.89	6.09	0.72	9	0.88	3	0.7	8.45
SHILL_e	0.3	0.16	0.67	0.99	0.97	0.93	5.52	0.66	6	0.55	7	0.83	8
agent^r	0.2	0.15	0	0.82	0.83	0.46	2.87	0.79	3	0.84	3	0.67	7.89
agent^e	0.3	0.14	0	0.67	0.66	0.44	2.96	0.86	3	0.82	3	0.77	7.55
agent^l	0.2	0.27	0	0.78	0.85	0.53	3.03	0.86	3	0.54	3	0.67	7.39
agent^f	0.3	0.13	0	0.48	0.68	0.48	2.83	0.79	3	0.79	3	0.77	7.01
agent^d	0.4	0.16	0.75	0.17	0.91	0.06	4.52	0.83	3	0.83	3	0.77	6.88
agent^o	0.3	0.29	0	0.7	0.66	0.61	3.28	0.66	6	0.55	7	0.83	6.79
agent^j	0.1	0.63	0	0.88	0.76	0.77	3.17	0.17	5	0.48	6	1	6.59
agent^m	0.1	0.5	0	0.93	0.65	0.88	3.09	0.17	5	0.5	6	1	6.5
agent^q	0.5	0.29	0.8	0.88	0.51	0.77	6.09	1	7	0	8	0.8	6.38
agent^t	0.2	0.25	0	0.93	0.71	0.77	3.23	0.52	8	0	5	1	6.3
agent^p	0.2	0.13	0.5	0.79	0.48	0.48	3.65	0.41	2	0.47	5	1	6.3
agent^c	0.1	0.1	0	0.86	0.73	0.38	2.29	0.21	5	0.18	6	1	5.95
agent^i	0.2	0.17	0.5	0.47	0.12	0.34	2.95	0.41	2	0.57	5	1	5.14
agent^k	0.2	0.17	0.5	0.49	0.63	0.42	3.5	0.34	2	0.57	2	0.5	5.05
agent^h	0.2	0.3	0.5	0.88	0.24	0.64	3.83	0.38	2	0.39	2	0.5	4.77
agent^g	0	0	0	0	0	0	0	0	0	1	0	1	4
agent^s	0	0	0	0	0	0	0	0	0	1	0	1	4
agent^h	0	0	0	0	0	0	0	0	0	1	0	1	4
agent^a	0.1	0.25	0	0.83	0.67	0.81	2.74	0.24	1	0	1	0	3.48
agent^n	0	0	0	0	0	0	0	0	0	0	0	0	0

Fig. 8. Example collusion scores and shill information for the hybrid strategy

Once a suspicious group is discovered, they are then further examined using a binding factor that indicates how similar the group members are based on aspects of their shill score ratings. The collusion score combines all of these ratings and gives each bidder a score based on the likelihood that they are engaging in collusive shilling behaviour.

The ability of the collusion score to detect collusive shill bidding was tested on simulated auction data. Experimental results indicate that colluding shills are able to decrease their shill scores, compared to a single shill. The collusion algorithm was able to successfully identify shills that engaged in the alternating bid strategy. The algorithm highlighted major colluding groups using the alternating auction strategy. When shills used a hybrid strategy, the algorithm partitioned potential colluding groups. Some innocent bidders may be incriminated and some shills may not be detected. However, this work forms the basis for more sophisticated shill detection techniques that minimise the extent of false incrimination and non-detection. Other future work involves examining the effects of multiple

seller collusion on shill detection.

REFERENCES

- [1] J. Schwartz and J. Dobrzynski, "3 men are charged with fraud in 1,100 art auctions on eBay," in *The New York Times*, 2002.
- [2] H. Shah, N. Joshi and P. Wurman "Mining for Bidding Strategies on eBay," in *SIGKDD'2002 Workshop on Web Mining for Usage Patterns and User Profiles*, 2002.
- [3] S. Rubin, M. Christodorescu, V. Ganapathy, J. Giffin, L. Kouger and H. Wang, "An Auctioning Reputation System Based on Anomaly Detection," in the *Proceedings of the 12th ACM Conference on Computer and Communications Security (CCS)*, pages 270–279, Alexandria Virginia, 2005;
- [4] J. Trevathan and W. Read "Detecting Shill Bidding in Online English Auctions," *Technical Report*, 2005. Available at <http://auction.maths.jcu.edu.au/research/shill.pdf>
- [5] J. Trevathan and W. Read, RAS: a system for supporting research in online auctions. *ACM Crossroads*, ed. 12.4, pages 23–30, 2006.
- [6] J. Trevathan and W. Read "Fraudulent and Undesirable Behaviour in Online Auctions," in the *Proceedings of the International Conference on Security and Cryptography (SECRYPT)*, pages 450–458, 2006.
- [7] J. Trevathan and W. Read, A simple shill bidding agent, in the *Proceedings of the 4th International Conference on Information Technology - New Generations*, (to appear), 2007. Available at <http://auction.maths.jcu.edu.au/research/agent.pdf>

Chapter 6

Software Bidding Agent Security

This chapter addresses automated bidding agent security. Previous literature on agent based negotiation is mainly concerned with the performance aspects of bidding agents. Security issues are largely neglected. Malicious bidding agents were created to evaluate the effectiveness of the proposed security mechanisms. Three papers are presented. The first paper describes the characteristics and operation of a basic shill bidding agent (see Section 6.1). The second paper introduces a novel approach for determining the extrema in a given data set in real-time, which is referred to as the Extremum Consistency (EC) algorithm (see Section 6.2). The third paper presents an adaptive shill bidding agent that uses the EC algorithm, to predict and refine its bidding behaviour over a series of auctions (see Section 6.3).

6.1 A Simple Shill Bidding Agent

This paper presents a software bidding agent that follows a shill bidding strategy. The malicious bidding agent was constructed to aid in developing shill detection techniques. The agent incrementally increases an auction's price, forcing legitimate bidders to submit higher bids in order to win the item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction. The agent's ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. This is the first documented bidding agent that perpetrates auction fraud. The simple shill agent is used in the construction of the adaptive shill bidding agent (see Section 6.3).

"A Simple Shill Bidding Agent" [10] is to appear in the 4th International Conference on Information Technology - New Generations 2007 (ITNG). ITNG is an annual event focusing on state of the art technologies pertaining to digital information and communication. The applications of advanced information technology to such domains as astronomy, biology, education, geosciences, and health care are among topics of relevance to ITNG. Visionary ideas, theoretical and experimental results, as well as prototypes, designs, and tools that help the information readily flow to the user are of special interest. The conference features keynote speakers, the best student award, the PhD student forum and workshops/exhibits from industry, government and academia. Proceedings are published by the IEEE Computer Society. This paper is available online via citeseer ¹. ITNG 2007 is to be held in Las Vegas, U.S.A.

¹<http://citeseer.ist.psu.edu/>

Publication in ITNG was made possible through funding by the Graduate Research Scholarship (GRS) and the School of Maths, Physics and IT.

A Simple Shill Bidding Agent

Jarrod Trevathan
School of Maths, Physics and IT
James Cook University
Email: jarrod.trevathan@jcu.edu.au

Wayne Read
School of Maths, Physics and IT
James Cook University
Email: wayne.read@jcu.edu.au

Abstract—Shill bidding is where fake bids are introduced into an auction to drive up the final price for the seller, thereby defrauding legitimate bidders. Although shill bidding is strictly forbidden in online auctions such as eBay, it is still a major problem. This paper presents a software bidding agent that follows a shill bidding strategy. The malicious bidding agent was constructed to aid in developing shill detection techniques. The agent incrementally increases an auction’s price, forcing legitimate bidders to submit higher bids in order to win the item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction. The agent’s ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. This is the first documented bidding agent that perpetrates auction fraud. We do not condone the use of the agent outside the scope of this research.

I. INTRODUCTION

Online auction fraud is rampant and every year increasingly costs victims millions in stolen money and sanity. Shill bidding is the act of introducing spurious bids into an auction on the seller’s behalf, with the intent to artificially inflate the item’s price. Bidders who engage in shilling are referred to as ‘*shills*’. To win the item, a legitimate bidder must outbid a shill’s price. If the shill accidentally wins, then the item is re-sold in a subsequent auction. Shill bidding defrauds legitimate bidders as they are forced into paying significantly more for the item.

There are many auction types (e.g., Vickrey, Dutch, etc.). The most popular is the English auction. In an English auction, bidders outbid each other in an attempt to win an item. The winner is the bidder with the highest bid. English auctions are employed in online auctions such as those offered by eBay¹ and ubid². The English auction is particularly susceptible to shill bidding practices. Shill bidding is strictly forbidden by commercial online auctioneers, and is a prosecutable offence (see [6]). However, the online environment makes shilling easy as bidders are anonymous.

There is often much confusion regarding what constitutes shill behaviour. Bidding behaviour that might seem suspicious, could in fact turn out to be innocent. Furthermore, a shill can engage in what seems to be a limitless number of strategies. This makes it difficult to detect shill bidding. While the online auctioneers monitor their auctions for shilling, there is little academic material available on proven shill detection techniques.

The Research Auction Server (RAS)³ at James Cook University, is an online server for conducting research into security issues regarding online auctions (see [7]). We are developing methods to detect fraudulent bidding behaviour such as shilling (see [9]). Both real and simulated auctions are conducted to test the effectiveness of the detection methods. To aid in testing, we developed a software bidding agent that bids in a manner consistent with a shill.

A *software bidding agent* is a program that bids on a human bidder’s behalf. Agents follow a predetermined strategy, typically with the goal of winning an auction for the minimal amount. In an English auction, a bidding agent is permitted to outbid any bid until the bidding price exceeds a maximum amount specified by the human bidder. In auctions that can last days or weeks, bidding agents remove the need for a bidder to constantly observe an auction. A bidding agent monitors the auction proceedings for any price activity, and responds in accordance with its programmed strategy.

Numerous bidding agents have been proposed in literature (see [1], [3], [4], [5]). The Trading Agent Competition (TAC)⁴ (see [10]) pits bidding agents against each other in an elaborate economic game. The TAC server allows bidders to write their own agents using an application programming interface. Agents are assessed on their ability to acquire resources in the most efficient manner. TAC is mainly concerned with furthering the performance aspects of (non-fraudulent) bidding agents. It does not focus exclusively on agent security. All TAC agents are required to behave strictly within the auction’s rules, and are disqualified if they act with a malicious intent to influence the auction proceedings.

To our knowledge no literature exists for a type of bidding agent we refer to as a *malicious bidding agent*. Similar to a virus or worm, a malicious bidding agent is a bidding agent that behaves with an intent to do an auction harm in some manner. This might be in the form of inflating the final price by shilling, attacking the cryptographic protocols of a “secure” auction system, or launching a denial of service attack against the Auctioneer.

Automated agents for conducting electronic commerce are becoming more common. The idea has been touted that all human input in auctions will eventually be done using autonomous bidding agents. However, such an environment

¹<http://www.ebay.com>

²<http://www.ubid.com>

³<http://auction.math.jcu.edu.au>

⁴<http://www.sics.se/tac/>

would definitely spawn undesirable behaviour (e.g., cheating, stealing payments, etc.). (See [8].) Furthermore, undesirable groups such as terrorists can obtain funds through fraudulent activities in auctions. In an extreme scenario, a malicious agent could hinder the world's stock exchanges, in an attempt to undermine the financial system. Therefore, the threat posed by malicious bidding agents to electronic commerce is very serious. As a result, no one is willing to totally trust agent-based negotiation.

Existing agents operate in a controlled and near perfect environment. As it is not permitted to use malicious agents in commercial online auctions, or in TAC, we have created an agent interface for RAS. RAS allows the agents to be tested in a controlled (and legal) manner.

This paper presents a software bidding agent that follows a shill bidding strategy. The agent incrementally increases the auction's price, forcing legitimate bidders to submit higher bids in order to win. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction. The agent's performance has been tested against non-fraudulent bidding agents in a simulated market place. Experimental results show that the agent is able to successfully increase the average winning price. This paper *does not condone* the use of these agents in any manner outside the scope of this research. By developing malicious bidding agents, we hope to better understand the characteristics of such agents, and how to protect against the damage they inflict.

This paper is organised as follows: Section II describes general shill behaviour. Section III presents a shill bidding agent that shills in a single auction. Section IV evaluates the agent's ability to shill, contrasting safe and risky approaches. Section V provides some concluding remarks and avenues for future work.

II. SHILL BEHAVIOUR

This section provides an insight into general shill behaviour. It describes a shill's characteristics and strategies, and contrasts shilling with another type of bidding behaviour referred to as *sniping*.

The main goal for shilling is to artificially inflate the price for the seller beyond what legitimate bidders would otherwise require to win the item. The pay-off for the seller is the difference between the final price and the uninflated price. A shill's goal is to lose each auction. A shill is not constrained by a budget, but rather a profit margin. If the shill wins, the item is resold in a subsequent auction. However, there is a limit on how many times this can be done. For each auction a shill wins, the seller incurs auction listing fees and is required to invest more time. Continual wins erode the profit from shilling on the item.

The shill faces a dilemma for each bid they submit. Increasing a bid could marginally increase the revenue for the seller. However, raising the price might also result in failure if it is not outbid before the auction terminates. The shill must decide whether to 'take the deal', or attempt to increase the pay-off.

On the contrary, a bidder's goal is to win. A bidder has a finite budget and is after the lowest price possible. Increasing a bid for a legitimate bidder decreases the money saved, but increases the likelihood of winning. The following outlines typical shill behaviour and characteristics:

- 1) A shill tends to bid exclusively in auctions only held by one (or a few) particular seller(s).
- 2) A shill generally has a high bid frequency. An aggressive shill will continually outbid legitimate bids to inflate the final price, until the seller's expected pay-off for shilling has been reached, or if the shill risks winning the auction (e.g., near the termination time or during slow bidding).
- 3) A shill has few or no winnings for the auctions participated in.
- 4) It is advantageous for a shill to bid within a small time period after a legitimate bid. Generally a shill wants to give legitimate bidders as much time as possible to submit a new bid before the auction's closing time.
- 5) A shill usually bids the minimum amount required to outbid a legitimate bidder. If the shill bids an amount that is much higher than the current highest bid, it is unlikely that a legitimate bidder will submit any more bids and the shill will win the auction.
- 6) A shill's goal is to try and stimulate bidding. As a result, a shill will tend to bid more near the auction's beginning. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the auction's end is risky as the shill could accidentally win.

Bid sniping is a type of bidding behaviour that is typically deemed undesirable. A bidder using a sniping strategy will only bid in the auction's closing seconds. The goal is to prevent other bidders from outbidding the sniper's bid, as they do not have time to respond. Sniping behaviour is shilling's exact opposite. A sniper's goal is to win the auction for the lowest price, whereas a shill's goal is not to win and to create the highest price. Sniping is often used as a preventative measure against shilling. A sniper cannot prevent shilling occurring during an auction. However, the sniper can prevent itself from being shilled. Unlike shilling, sniping is permitted on eBay although its use is discouraged.

III. A SHILL BIDDING AGENT

This section presents a shill bidding agent that uses bogus bids to inflate an auction's price. We describe the agent's goals and strategic directives. Each directive plugs into the agent interface, which dictates its bidding behaviour.

A. Components of an English Auction

In order to participate in an auction, a bidder must *register*. They are provided with a unique bidder id, b_{id} , which they use to submit bids. During the *initialisation* stage, the Auctioneer sets up the auction and advertises it (i.e., item description, starting time, etc). An auction is given a unique number, a_{id} , for identification purposes. In the *bidding* stage, a bidder computes his/her bid and submits it to the Auctioneer. The

```

A - Directive D1
float D1(current price p) {
    return minimum price
}

B - Directive D3
boolean D3(risk limit θ) {
    calculate d
    calculate tS
    if (tc ≤ tS) then
        return true
    else
        return false
}

```

Fig. 1. Pseudocode for Shill Bidding Agent Directives 1 and 3

agent can place a bid in auction a_{id} , for price p' , by invoking the **submit bid**(a_{id} , p') function.

The Auctioneer must *supply intermediate information* to the agent pertinent to the auction's current state. The agent can request a price quote for a particular auction by invoking the **obtain price quote**(a_{id}) function. This includes the start, end and current time for the auction, and the starting bid (if one exists). It is assumed that the agent has access to the entire bid history up to the current time in the auction. The history can be considered as an ordered set $\mathcal{H} = \{h_1, h_2, \dots, h_n\}$, $|\mathcal{H}| = n$, that contains price quote triples $h_i = (time, price, b_{id})$, where $1 \leq i \leq n$. The last element is the latest price quote for the auction (i.e., h_n is the current highest bid).

Finally, during the *winner determination* stage, the Auctioneer chooses the winner according to the auction rules (e.g., who has the highest bid, whether the reserve has been met, etc.).

B. Shill Agent Directives

The agent's goal is to maximise the profit from shilling, while avoiding winning the auction. A shill that wins the auction is deemed to have failed. We propose a set of directives that a shill must adhere to. If these directives are satisfied, the agent submits a bid. Each directive is described in turn.

D₁ - Bid minimum amount required. As part of the price quote, the Auctioneer also provides the minimum amount required to outbid the highest bid. This is usually calculated as a percentage of the current high bid, or determined according to a scalable amount depending on the current high bid. $D_1(p)$ is a function that takes the current price, p , and returns the minimum amount the shill should bid. (See Figure 1A.)

D₂ - Bid quickly after a rival bid. The agent must bid immediately in order to influence the other bidders for the maximum time. The agent's location affects its response speed to rival bids. There are two main placement options:

- *On the Auctioneer.* The agent can respond instantly to a rival bid;
- *As a client of the Auctioneer.* The agent must periodically poll the Auctioneer to check if a new rival bid has been submitted. The length of the polling interval limits shilling.

The first approach uses the Auctioneer's computer to run the agent program. For example, eBay's proxy bidding system functions in this manner. However, this places a drain on the Auctioneer's computational and storage resources. As in TAC, the second approach allows bidders to host the agent on their own computer and interact with the Auctioneer via an application programming interface. This distributes the computational burden among the bidders. However, this requires a permanent connection to the Auction server and results in communication overhead. Furthermore, network delays and security threats can influence the agent's operating speed and integrity. RAS facilitates both types of agent placement.

D₃ - Don't bid too close to the auction's end. If the shill bids too close to the auction's end, it risks winning. To avoid this, the agent has a risk limit, θ , $0 \leq \theta \leq 1$. The agent is prohibited from bidding if the auction is more than $(\theta \times 100)\%$ complete. We refer to this as *shill time limit*, and denote it as t_S . t_S is the absolute bound after which a shill is prevented from bidding. Let t_0 and t_e be the starting and ending times respectively for an auction. The auction's duration, d , is calculated as $d = t_e - t_0$. The shill time limit, t_S is calculated as $t_S = \theta d$. Let t_c be the current time in an auction. The agent can only submit bids when $t_c \leq t_S$. Larger values of θ increase the risk that a shill might win an auction. $D_3(\theta)$ is a function that takes the risk limit θ , and returns true or false regarding whether the agent should continue to bid. (See Figure 1B.)

D₄ - Bid until the target price has been reached. The reserve price can influence a bidder's strategy. It has been argued that auctions with lower reserve prices attract more bids, whereas bidders are deterred by high reserve prices. There are three factors that influence an auction's final price: the *reserve price*, r ; the *seller's true valuation*, sv ; and a *bidder's true valuation*, bv .

If the seller lists r below sv , s/he risks selling the item at a loss. This is shown in Figure 2A. If the seller lists r above sv , this may potentially increase the seller's profit (see Figure 2B). However, this may also deter bidders if it is above bv (see Figure 2C), in which case the item is not sold, or if there is no reserve, it is sold for a loss. Strategies regarding the choice of reserve price, touches on the complex topic of




Fig. 2. Pricing and Strategies

reserve price shilling, which is beyond the scope of this paper. For simplicity, we assume that the seller's best strategy is to choose $sv = r$ (see Figure 2D). In this case, it is clear to see that profit occurs where $bv > r$ (see Figure 2E).

The shill agent's primary goal is to inflate the price to at least r , (and by assumption sv). Its secondary goal is to etch out further profit in the case where $bv > r$. In order to know when to stop bidding, the shill is given a *target price*, which is denoted by α . The higher α is, the more profit the shill can attain. However, this increases the chances of failure (i.e., winning the auction). Ideally, the choice of α should be $r \leq \alpha \leq bv$ (see Figure 2F). When choosing α , the shill must ensure that it is greater than or equal to r . Otherwise, the item might be sold at a loss. If $bv = r$, then the seller can be sure that shilling will inflate the price to at least to r (see Figure 2G). bv must be greater than r , for the shill agent to be able to etch out profit. Profit in this case is measured as the difference between the inflated price and r (see Figure 2H).

Dumas *et al* [2] performed an analysis of datasets from eBay and Yahoo⁵ and showed that the final prices of a set of auctions for a given item are likely to follow a normal distribution. This is due to a given item having a more or less well-known value, around which most of the auctions should finish. Figure 3 illustrates a typical closing distribution. bv can typically be deemed to be at the distribution's centre. However, this should only be used as a short to medium term indicator, as permanent price changes over time do occur.

$D_4(p, \alpha)$ is a function that takes the current price p , the shill's target price α , and returns true or false regarding

⁵www.yahoo.com/auction




Fig. 3. Final price bid distribution over a series of relatively concurrent auctions

whether the agent should continue to bid. The agent will only bid when the current price p , is less than or equal to α . (See Figure 4A.)

D_5 - Only bid when the current bidding volume is high.

The agent should preferably bid more towards the auction's beginning and slow down towards the shill time limit, unless the bidding activity is high. That is, the bidding volume must increase throughout the auction for the shill to maintain the same bid frequency. The agent uses the bid history \mathcal{H} , to analyse the current bid volume and decide whether to submit a bid. The agent observes the previous number of bids for a time interval. If the number of bids for the period is below a threshold, then the agent does not submit a bid.

First, we must determine each bid's normalised time in terms of the current time, t_c . This is represented as δh_i , $1 \leq i \leq n$. Let t_i be the time for h_i . $\delta h_i = \frac{t_i - t_0}{t_c - t_0}$, where $0 \leq \delta h_i \leq 1$. The risk value μ , $0 \leq \mu \leq 1$, represents how far back in history from the current time that the agent will observe. For example, if $\mu = 0.2$, then the agent will only look for bids that were submitted in the final 20% of the normalised time period of the auction thus far (i.e., $\delta h_i \geq (1 - \mu)$). κ denotes the number of bids submitted in the last $(\mu \times 100)\%$ of elapsed time from t_c . Increasing μ , increases the level of risk for the shill (as the agent is influenced by bids further in the past). κ is calculated as follows:

$$\kappa = \sum_{i=1}^{|\mathcal{H}|} j \text{ where } j = \begin{cases} 1, & \text{if } \delta h_i \geq (1 - \mu) \\ 0, & \text{otherwise} \end{cases}$$

where $0 \leq \kappa \leq |\mathcal{H}|$.

Next, κ must be weighted depending on the current time in relation to the shill time limit, t_S . An increase in the trading volume is required towards t_S , in order for the agent to continue bidding at the same rate it did earlier in the auction. The normalised current time δt_c , in relation to t_S is calculated as $\delta t_c = \frac{t_c - t_0}{t_S - t_0}$, where $0 \leq \delta t_c \leq 1$. If $\delta t_c < 0.5$, then the agent will only require one bid to be submitted by a rival before it bids. When $\delta t_c \geq 0.5$, the agent requires at least two bids before it will submit a bid. This ensures that later in the auction, the agent will only respond to more aggressive rival behaviour.

When no bids have been submitted for an auction, the shill agent will attempt to stimulate bidding by submitting the first


```

A - Directive D4
boolean D4(current price  $p$ , target price  $\alpha$ ) {
    if ( $p \leq \alpha$ ) then
        return true
    else
        return false
}

B - Directive D5
boolean D5(risk limit  $\mu$ ) {
    if ( $|\mathcal{H}| = 0$ ) then
        return true
    calculate  $\delta h_i$  for all  $i$ ,  $1 \leq i \leq n$ 
    calculate  $\kappa$ 
    calculate  $\delta t_c$ 
    if ( $\kappa > 0$ ) then
        if ( $\delta t_c < 0.5$ )
            return true
        else if ( $\kappa > 1$ ) then
            return true
    return false
}

```

Fig. 4. Pseudocode for Shill Bidding Agent Directives 4 and 5

bid. This is a common practice in auctions. It is intended that psychologically the presence of an initial bid raises the item's worth, as competitors see that it is in demand. $D_5(\mu)$ is a function that takes the risk limit μ , and returns true or false regarding whether the agent should continue to bid. (See Figure 4B.)

The aforementioned directives govern the agent's operation depending on the state of the auction. The following pseudocode illustrates the agent's behaviour:

```

shill agent( $aid$ ,  $\alpha$ ,  $\theta$ ,  $\mu$ ) {
    do {
        obtain price quote( $aid$ )
        if ( $D_3(\theta)$  AND  $D_4(p, \alpha)$  AND  $D_5(\mu)$ )
            submit bid( $aid$ ,  $D_1(p)$ )
    } while ( $D_3(\theta)$  AND  $D_4(p, \alpha)$ )
}

```

The agent initially requests a price quote. If no bids have been submitted, then D_5 returns true and the agent submits a bid for the amount returned by D_1 . The agent then repeatedly requests price quotes to ensure that it is able to bid quickly if there is a rival bid (i.e., D_2). When a rival bid is submitted, the agent will bid only if the remaining directives, D_3 , D_4 , and D_5 are satisfied. The agent executes in this manner (requesting and evaluating price quotes), until either D_3 or D_4 becomes false.

IV. PERFORMANCE

The agent was implemented on RAS and has been tested with other types of bidding agents in a simulated auction market. The shill agent was assessed on its ability to inflate an auction's final price. The agent was considered *successful* if the winning price of a rival bid equals or exceeds α . The agent was considered *unsuccessful* if it won the auction, or if it failed to inflate the final price to α prior to ceasing bidding.

The shill agent was pitted against four Zero Intelligence agents (ZI) (see [4]). ZI agents are designed to simulate an ordinary bidder in an auction. Each ZI agent is assigned a

random amount between \$0.05 and \$10.00 (according to a uniform distribution), which it submits as a proxy bid at a random time during the auction. It was assumed that there was no reserve price.

Tests indicated that large numbers of ZI agents makes the shill less effective at influencing the final price (as there is no need to stimulate bidding). The average final price and bid volume for an auction **without shilling** were \$5.90 and 3.95 bids respectively. The standard deviation for closing prices was \$1.95.

Claim 1 *The shill agent raised the average winning price compared to auctions it didn't participate in.* The average price was 1% – 25% higher in auctions where shilling occurred (depending on the shill's risk profile). Where $\alpha < bv$, the shill can influence the price the most. After this, the average final price becomes affected by the shill's winning bids. This can be seen in Figure 5 A.

Claim 2 *The bidding volume increased in auctions where the shill agent was present.* Introducing a shill agent increased the average bid volume by up to 420%. Much of this can be attributed to the agent incrementally outbidding proxy bids. In general, the higher α , the more bids a shill will have to submit. This can be seen in Figure 5 B.

Claim 3 *Riskier shill agents acquire more profit, but at the expense of an increased number of failures.* We conducted numerous tests that altered the shill's risk parameters, α , θ and μ . More risk adverse shills tended to fail by not meeting α , whereas riskier shills tended to fail by winning the auction. Figure 5 C shows how the agent's success rate decreases with increases in α . Figure 5 D shows how the percentage of these failures that are due to winning the auction.

Claim 4 *A shill agent that places a single proxy bid for the target price α , is less effective.* Tests showed that shills employing this strategy tend to fail more by winning the




Fig. 5. Graph illustrating the agent’s performance with increasing risk factors θ and μ . The horizontal axis represents the target price α . **A** shows the increase in average final price with increases in α . Likewise **B** shows the increase in average number of bids per auction with increases in α . **C** illustrates how the agent’s success rate decreases with α and **D** shows the increase in failures due to winning the auction.

auction. An agent employing the shilling strategy outlined in this paper uses θ and μ to help determine whether it is safe to continue. Psychologically, the use of a larger number of small bids is also more likely to lure bidders into placing higher bids. Furthermore, it is less suspicious for the shill to slowly inflate the price, rather than enter an initial large amount.

Claim 5 *The shill agent was less effective against sniping agents.* The shill agent was pitted against varying numbers of ZI agents following a sniping strategy (see [5]). As the number of sniping agents reached saturation (i.e., one shill vs all sniping agents), the shill’s ability to influence the price decreased to 0. Thus the shill failure rate was 100% due to α not being met (unless normal bidding inadvertently reached α). Bid volume was also significantly lower.

V. CONCLUSIONS

This paper presents a malicious software bidding agent that follows a shill bidding strategy. The agent incrementally increases an auction’s price forcing legitimate bidders to submit higher bids in order to win the item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction.

The shill bidding agent has been implemented on RAS to aid us in developing shill detection techniques. The agent’s ability to shill was tested using a simulated auction market involving other bidding agents. The agent raised the overall bidding volume as well as increased the average final price across all auctions. Shills with a higher profit risk factor managed to acquire a larger amount of profit. However, this came at the expense of an increased number of failed auctions. Tests

conducted showed that the shilling strategy employed by the agent, is superior to placing a single proxy bid for the target price. Engaging in sniping reduces the effectiveness of a shill bidder.

In an extended version of this paper, we have developed an adaptive shill bidding agent. When used on a series of auctions for substitutable items, the agent is able to revise α and other risk factors based on its past performance. In future work we plan to investigate shill agents which engage in collusive behaviour to avoid detection.

REFERENCES

- [1] Cliff, D., Minimal-intelligence agents for bargaining behaviours in market-based environments, *Hewlett Packard Labs, Technical Report HPL-97-91*, 1997.
- [2] Dumas, M., Aldred, L. and Governatori, G., A probabilistic approach to automated bidding in alternative auctions, In *Proceedings of the eleventh international conference on World Wide Web*, 99–108. ACM Press, 2002.
- [3] Gjerstad, S. and Dickhaut, J., Price formation in double auctions, *Games and Economic Behavior*, 22, 1–29, 1998.
- [4] Gode, D. and Sunder, S., Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality, *Journal of Political Economy*, 101, 119–137, 1993.
- [5] Rust, J., Miller, J. and Palmer, R., Behaviour of trading automata in a computerized double auction market, In *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1992.
- [6] Schwartz, J. and Dobrzynski, J., 3 men are charged with fraud in 1 100 art auctions on eBay, *The New York Times*, 2002.
- [7] Trevathan, J. and Read, W., RAS: a system for supporting research in online auctions, *ACM Crossroads*, 12.4, 23–30, 2006.
- [8] Trevathan, J. and Read, W., Undesirable and fraudulent behaviour in online auctions, In *Security and Cryptography Conference (SECURITY)*, 450–458, 2006.
- [9] Trevathan, J. and Read, W., Detecting shill bidding in online English auctions, *Technical Report*, 2006. Available at: <http://auction.maths.jcu.edu.au/research/shill.pdf>
- [10] Wellman, M. and Wurman, P., The 2001 trading agent competition, *Electronic Markets*, 13.1, 4–12, 2003.

6.2 A New Approach to Avoiding the Local Extrema Trap

This paper introduces the Extremum Consistency (EC) algorithm for avoiding local maxima and minima in a specialised domain. The most notable difference between this approach and others in the literature is that it places a greater importance on the *width* or *consistency* of extremum than on height or depth (amplitude). Short-term, high amplitude extrema can be encountered in many typical situations (such as noisy environments or due to hardware inaccuracies) and can cause problems with system accuracy. The EC algorithm is far less susceptible to these situations than hill climbing, convolution, thresholding etc., and tends to produce higher quality results. This paper details the algorithm itself as well as presenting the results of practical experimentation illustrating the superiority of EC over other forms of local extrema avoidance in three large-scale practical applications. The EC algorithm was originally constructed for use in handwritten signature verification (refer to Trevathan and McCabe 2005 [3]). The adaptive agent applies the EC algorithm to predict maximum and minimum prices in an auction dataset (see Section 6.3).

“*A New Approach to Avoiding the Local Extrema Trap*” [11] was presented at the 13th Biennial Computational Techniques and Applications Conference - CTAC 2006, and has been accepted for publication in the ANZIAM journal. CTAC is organised by the special interest group in computational techniques and applications of ANZIAM, the Australian and New Zealand Industrial & Applied Mathematics Division of the Australian Mathematical Society. CTAC provides an interactive forum for researchers interested in the development and use of computational methods applied to engineering, scientific and other problems. The CTAC meetings have been taking place biennially since 1981. CTAC 2006 is the 13th meeting in the series and is the first to be held in North Queensland. CTAC 2006 was hosted by the School of Mathematical and Physical Sciences, James Cook University. Papers were reviewed by a scientific committee and additional conference attendees. This paper is available online via citeseer ².

“*Preprocessing of On-line Signals in a Noisy Environment*” [12] is to appear in the 4th International Conference on Information Technology - New Generations 2007 (ITNG). ITNG is an annual event focusing on state of the art technologies pertaining to digital information and communication. The applications of advanced information technology to such domains as astronomy, biology, education, geosciences, and health care are among topics of relevance to ITNG. Visionary ideas, theoretical and experimental results, as well as prototypes, designs, and tools that help the information readily flow to the user are of special interest. The conference features keynote speakers, the best student award, the PhD student forum and workshops/exhibits from industry, government and academia. Proceedings are published by the IEEE Computer Society. This paper is available online via citeseer ³. ITNG 2007 is to be held in Las Vegas, U.S.A.

I am part of the organising and scientific committee for CTAC 2006. These papers are co-authored with Dr Alan McCabe of Fern Computer Services in Belfast, Ireland. Dr McCabe’s research interests include Artificial Intelligence and Biometric Security. I was responsible for presenting the paper at the conference and guiding it to publication. This involved the majority of the editorial tasks. I am solely responsible for creating and presenting the poster. Not documented in the paper is a complexity analysis and additional notes regarding the algorithm. Publication in ITNG was made possible through funding by the Graduate Research Scholarship (GRS) and the School of Maths, Physics and IT.

²<http://citeseer.ist.psu.edu/>

³<http://citeseer.ist.psu.edu/>

A New Approach to Avoiding the Local Extrema Trap

Alan McCabe
Fern Computer Services Ltd.
Belfast, Ireland
Email: alan@mymait.com

Jarrod Trevathan
School of Maths, Physics and IT
James Cook University
Email: jarrod.trevathan@jcu.edu.au

Abstract—This paper introduces the Extremum Consistency (EC) algorithm for avoiding local maxima and minima in a specialised domain. The most notable difference between this approach and others in the literature is that it places a greater importance on the *width* or *consistency* of an extremum than on its height or depth (amplitude). Short-term, high amplitude extrema can be encountered in many typical situations (such as noisy environments or due to hardware inaccuracies) and can cause problems with system accuracy. The EC algorithm is far less susceptible to these situations than hill climbing, convolution, thresholding etc., and tends to produce higher quality results. This paper describes the algorithm and presents results from practical experimentation, which illustrates its superiority over other forms of local extrema avoidance in three real-world applications.

I. INTRODUCTION

This paper presents a new algorithm for avoiding local maxima and minima in specialised environments. The embedding of this algorithm in a number of practical systems has resulted in significant improvements in accuracy when compared with the use of other classical local extremum avoidance algorithms.

Note that throughout this paper the discussion will tend to focus on minima (also known as valleys or troughs) in the interests of brevity. Discussions can trivially be adapted to avoidance of maxima or peaks. Additionally, chiefly for convenience, this algorithm has been given a name: Extremum Consistency or EC. The meaning behind the name will be made clear later in the paper.

The underlying problem lies in deciding whether a given extremum represents a “true” extremum. This of course is not a new problem to computer science, but it has never been approached in the particular way described here.

There are several existing algorithms for avoiding local extrema, however none have proven to be sufficiently effective in the specific domain outlined in Section II. The most appropriate of these existing algorithms were implemented in complete systems and the results of these are presented below.

The paper is organised as follows, Section II describes the problem domain and motivation for the algorithm, Section III presents the algorithm and Section IV outlines some of the specific software applications where the algorithm has been successfully applied. Section V presents the concluding remarks and directions for future work.

II. THE PROBLEM DOMAIN - MOTIVATION

This problem domain is based on iterative improvement strategies that attempt to find maxima or minima while “traversing” or “stepping” along a certain stream.

Abstractly, what is required in this domain is for an algorithm to move in the direction of decreasing value until a minimum is reached. Once there, the minimum’s location is recorded and, depending on the application, it either proceeds in the opposite direction looking for a corresponding maximum, or “jumps” out of the minimum and begins the process again. This search continues until all of the minima (and maxima if required) are mapped and then processing of these points is performed (see Section IV for more information on the specific applications). Because of this pattern of searching, any false (in a sense, local) minimum encountered will adversely affect the performance of the system, not only because the recorded position is incorrect, but starting the search for the subsequent extremum too early may propagate that error.

The main problem therefore lies in detecting the *true* (or global) extrema in a volatile and often noisy environment. The EC algorithm is a specific approach to deciding whether a particular extremum represents a true extremum. Figure 1 pictorially shows some of the situations typically encountered in this environment. The horizontal-axis in these diagrams represents time and the vertical-axis can represent various observations such as velocity, direction or temperature.

The initial motivation for this algorithm came when developing signature verification software in [2]. This approach was based on detecting the order of “turning points” (in other words, minima and maxima) in the pen tip direction and processing the locations of those turning points. The EC algorithm was implemented in this application with significant improvements in system accuracy. See Section IV for more information on this and other application areas.

The challenge for this and similar algorithms is to ignore meaningless small fluctuations that appear in the stream, while recording the *meaningful* fluctuations. The problem lies in distinguishing between the two. Local extrema can enter the data as a result of various aspects, such as quantisation noise, rounding problems (discussed below) or, in handwriting based applications (discussed in Section IV), something as simple




Fig. 1. This figure represents some local minima situations that are typically encountered in processing the input stream. The horizontal-axis represents increasing time and the vertical-axis can represent various stream types such as velocity, direction or temperature. Specifically, (a) contains a valid minimum, (b) contains only a single valid minimum (the appearance of two minima is caused by some noise in the stream, so the second should be ignored) and no others contain any “true” minima, according to our definition (the minima in (c) and (f) are probably a result of noise in this environment, rather than genuine minima, and the same could be said for (d) and (e)). An effective algorithm should reflect this.

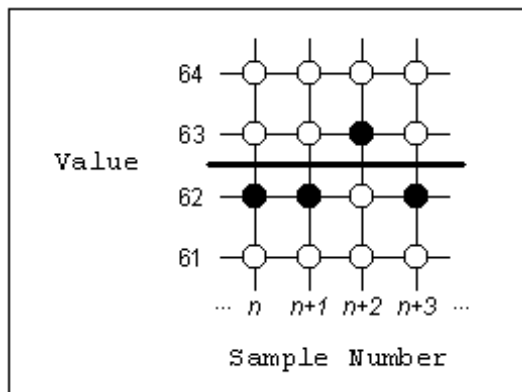


Fig. 2. Situations like this are the result of the discrete resolution of the hardware used to capture a stream. The bold horizontal line represents the actual position of the data and the black dots represent the recorded values. Time is represented on the horizontal-axis. This situation typically arises when the hardware is a graphics tablet or pen-based computer that rounds the position of the pen tip to the nearest pixel, but also arises with (say) temperature observations when the actual temperature is rounded to the nearest tenth of a degree for recording.

as shaky hands or the writing surface itself. These local extrema can be of varying scale, making them more difficult for conventional algorithms to overcome.

The aforementioned rounding problem occurs due to the discrete resolution of hardware such as graphics tablets or pen-based computers used to capture handwriting. The problem occurs when the actual handwriting path travels directly between two neighbouring pixels. The tablet must “round” the pen tip location to the nearest pixel. Occasionally, slight variations in pen tip pressure cause the rounding to be done to a different pixel. Figure 2 shows an example of this. The resulting handwriting path then looks (to the system) like that shown in Figure 1(b) or Figure 1(c).

III. THE ALGORITHM

The algorithm is compact, requiring very little stored data (four integers), no search tree, and is implemented via a series of comparisons done while traversing the surface of the feature space. Additionally, the algorithm is only executed




Fig. 4. An illustration of step and width calculation. Valid steps occur between time points 0 and 1, 1 and 2, 3 and 4 and 5 and 6. Backward steps occur between time points 4 and 5, 6 and 7, and between 9 and 10.

when a potential minimum is encountered so the effect on the efficiency of the overall system is slight.

The major difference between this algorithm and others is that it examines, primarily, the *width* (or perhaps more accurately the *consistency* or *duration*) of the minima. Most other algorithms (such as convolution and thresholding) place more emphasis on the *depth* of the minima. The use of the term *width* here differs slightly from an intuitive understanding of the width of a valley. The width of a valley is best explained by considering the initial valley downslope and the following upslope separately. It is defined as the number of “steps” encountered in its traversal where a step, in a downslope, refers to a decrease in height below the value of the current minimum. The more of these decreases there are, the larger the number of steps. Once these two values are found, the width of the valley is defined as the *minimum* of the individual width values.

Figure 3 presents a finite state machine illustrating the EC algorithm’s operation. The remainder of this section contains descriptions to accompany the illustrations.

Figure 4 illustrates an example of step and width calculation. Steps are defined as movements in the direction of a




Fig. 3. A finite state machine expressing the EC algorithm for finding the “width” of, or number of “steps” in, the initial downslope of a valley. The movements between vertices (states) are defined by the comparison between points in the input stream and the comparisons are included on the edges in the diagram. Additionally there are actions to be performed when some vertices are reached - these are also included in the diagram. Note that in this table, $h[i]$ refers to the i^{th} element in the list of stream values h . Once the `backSteps` parameter exceeds the `tolerance`, the algorithm enters the dead state D and we have the width of the slope in the `steps` parameter. The width of the upslope is similarly calculable, with the inversion of various state transition conditions and the width of the valley itself is then the minimum of the downslope width and upslope width. Likewise, the width of peaks can be found with minimal modifications to the algorithm.

particular extremum - for example, in Figure 4 the movement between $Time = 0$ and $Time = 1$, $Time = 1$ and $Time = 2$, $Time = 3$ and $Time = 4$ and $Time = 5$ and $Time = 6$ each constitute a single step. The term “backward steps” is now also defined as movements *away* from the extremum, beyond the current maximum. For example, in Figure 4, the movement between $Time = 4$ and $Time = 5$, $Time = 6$ and $Time = 7$ as well as between $Time = 9$ and $Time = 10$ can be thought of as taking a backward step. A tolerance parameter determines how many backward steps are accepted before the algorithm terminates.

Upon termination we have the value for the slope’s width. Tolerance then becomes a significant factor as it represents the amount of time spent looking for a “better” minimum before giving up and accepting the one we have. If the minimum’s location is the only information sought (as is often the case, depending on the application), the goal has been achieved and the algorithm ends. If, however, the width of the entire valley is sought (as in the application described in Section IV-A) then the search for the top of the post-valley upslope takes place, starting from the valley floor.

Recall that the width of the entire valley is taken as the minimum of the width values of the downslope and the upslope. It is important to take the minimum because otherwise a very large downslope with a very small upslope would erroneously appear as a very large valley. For example, see




Fig. 5. An illustration of a large downslope with a small upslope appears in the valley at $Time = 2$. In this case the valley itself is actually very shallow and using the maximum or mean slope size would make the valley appear erroneously large - the use of the minimum slope size is much more accurate. By placing a threshold on the width, this valley will ideally be ignored and that at $Time = 5$ would be taken as the more appropriate minimum.

Figure 5.

With the width of the valley obtained it is then just a matter of setting a threshold on which width sizes will be considered large enough to constitute a genuine valley (and similarly for peaks). The calculation of the optimal threshold was done during a training phase via brute force experimentation for

each application. Results typically worsened considerably for a threshold value of five or over and for all applications in Section IV a threshold value of just two proved to be globally optimal.

The final phase in experimentation involved taking both an extremum's height/depth and width value into account, in an attempt to obtain an even more accurate estimate of "true" extrema. The simplest and most successful combination method was to simply take the product of the valley width and depth (or height). The result of this was that deeper valleys were now considered "better" minima than shallow valleys with similar width, which intuitively seems a more desirable effect. In the experiments conducted, this approach consistently provided the best overall results.

Section IV describes a number of projects in which the EC algorithm has been successfully implemented. Subsection IV-A presents a detailed comparison of the accuracy and execution speed of this algorithm versus other approaches such as convolution.

IV. SUCCESSFUL APPLICATIONS

It is worth noting at this point that the algorithm presented in Section III is used essentially as a pre-processing filter, the output of which serves as input to another stage (for example, as part of a signature verification system). The only real way that success of the EC algorithm is measured is by evaluating the success of the resulting application as a whole. This section discusses some of the application areas in which the EC algorithm has been successfully employed.

A. Direction Based Handwritten Signature Verification

This was the first project to benefit from the application of the EC algorithm as it relies heavily on extrema detection. It involved the design and development of a dynamic signature verification system [2]. At its most basic level this system tracked the direction of the pen-tip when performing a signature. Peaks and valleys (maxima and minima) were detected in both the horizontal and vertical directions, ordered and converted into a character string (see [2] for more detail).

The initial attempt at valley detection was to implement a naive gradient-descent algorithm that simply traversed in the direction of non-increasing value until the point where a greater value was encountered. This previous (lowest) point was then deemed to be the minimum (note that the surface is only one-dimensional, so there is no choice as to which direction to take when traversing). The problem with this approach was that noise periodically appeared in the stream producing false extrema. Additionally there was a problem with the rounding of the pen-tip location to the nearest pixel in the hardware device which also produced false valleys and peaks (see Figure 2). This often resulted in the character string becoming somewhat mis-representative of the signature, degrading the whole system's effectiveness.

The overall error rate (sum of false rejection and false acceptance rates) for the first version of the system was 6.9%. This was improved dramatically to 2.9% with the

implementation of the EC algorithm to better detect the valleys and peaks. Specifically the false rejection rate (the proportion of genuine signatures rejected as forgeries) was improved from 4.3% to 0.9% and the false acceptance rate (forgeries accepted as genuine signatures) was improved from 2.6% to 2.0%. As mentioned in Section III, the final phase of EC involved combining the width of an extrema with its "amplitude" (height/depth). The purpose of this being to assign a higher importance to extrema with higher amplitude, versus lower amplitude extrema of similar width. The best combination method was to take the product of the two values, which resulted in a significant improvement of the overall error rate to 2.3%.

In other attempts to improve the error rates, two more algorithms were implemented: simple removal of small valleys or peaks (called "thresholding") and basic convolution of the data prior to gradient-descent/hill-climbing. Thresholding simply involved determining the distance (in pixels) between the peak's location and the preceding valley's location (that is, the depth of a valley or height of a peak). If this distance was below a specified threshold then that peak was ignored and the traversal continued in the same direction.

Examples of situations where thresholding was successful can be seen in Figure 1(b), (c) and possibly (d). However, the algorithm's overall performance (in terms of error rate) was quite poor. The reason it seems, is that a valley's *depth* alone, while obviously containing useful information, is not the best validity indicator (at least in this environment), but rather the consistency/duration is more important.

Simulated annealing is an approach which avoids local extrema by "jumping ahead" some random distance when an extremum is encountered, to try to find a "better" extremum [7]. While this is very effective in some domains, it was envisaged that simulated annealing would not work well in this environment. This was because there are often long periods in which the stream value remains the same (plateaus - see Figure 1(e)), which can vary greatly in their duration. Small jumps ahead will work on many occasions, but not in situations such as this. Large jumps ahead will work in many situations also, but will tend to jump over smaller details. In the interest of experimentation, a modified simulated annealing algorithm from [7] was implemented. The number and size of jumps were limited by threshold values, which were optimised through trial-and-error. The unsuitability of this approach was reflected in the poor error rate of 24.0%.

Convolution is considered one of the most useful method of "smoothing out" or "averaging" one-off "bumps" or random noise while attempting to preserve those extrema which are truly indicative of the pen-tip direction. The basic idea behind convolution is that a window of some finite length (convolution matrices are possible in environments of higher dimensionality) is scanned across the stream of values [1]. The output pixel is the weighted sum of the input pixels within the window where the weights can be adjusted to perform various filtering tasks - when smoothing is performed the weights are generally all equal.

After the input stream was convoluted the hill-climbing/gradient-descent approach was used to obtain the extrema. Multiple attempts were made with convolution using window sizes varying from one (the trivial case) up to fifty, with the optimal window size (that which resulted in the lowest error rate over the entire signature database) found to be five. Experiments were also conducted with the number of iterations of convolution performed, allowing for the possibility that the first convolution run didn't smooth out all of the irrelevant extrema and further iterations were necessary. The best overall results were obtained after a single iteration of convolution with the results progressively deteriorating with further iterations, indicating that some of the true extrema were being incorrectly smoothed out.

There was also some experimentation with combining convolution with the EC algorithm. This involved using convolution to smooth the input stream before applying EC to detect the extrema. This approach proved to be more successful than convolution with hill-climbing but less successful than EC alone. The reason for this is most probably related to convolution smoothing out small but meaningful extrema.

Table I summarises the error rates of the implemented approaches used in the signature verification system. The EC algorithm is clearly superior in this environment.

The other advantage of the EC algorithm over convolution is execution speed. In a real-time application like signature verification, execution speed can become a serious issue. In order to perform convolution an entire extra layer of computation is required, as convolution of the raw data must be done prior to obtaining the extrema, whereas with the EC algorithm the checking is done at the same time as the search for extrema. Additionally, convolution can become quite expensive using a large window or with a large raw data size (for example, a typical data size in the application described in Section IV-C is over 12,000 entries).

Empirical experimentation with the signature verification system has found that convolution causes an average slowdown of 20-25% (depending on system parameters). The number of extra calculations required in the EC algorithm compared to naive hill-climbing is almost negligible with the slowdown of

the signature verification system experimentally found to be less than 3%.

B. Velocity Based Handwritten Signature and Password Verification

The EC algorithm was also used in a handwritten password verification system [3]. The first step in this system (as well as many other handwriting based systems like character recognition algorithms [5]) was to segment the writing stream into its conceptually significant or constituent parts, commonly known as *strokes*. The strokes are continuous "pen-down" segments of writing bounded by consecutive minima in the pen-tip velocity. The approach then is to extract properties of these strokes and model these properties using a hidden Markov model or neural network.

An approach along these lines was presented previously in [3] and it also made successful use of the EC algorithm. The most naive method of obtaining the velocity minima is a basic gradient-descent algorithm. This was seen as an obvious application for the EC algorithm and it was implemented immediately. A simple gradient-descent implementation was also performed for comparative purposes. Simple gradient-descent produced a total error rate of 2.7% for password verification compared with 0.79% using EC with width only, and 0.64% when using EC with the product of width and height to do the segmentation. A similar EC implementation was also used as a method of signature segmentation in recent extensive studies [4], [8].

C. Physiology Research - Tracking Fluctuations in Infant Face Temperature

This physiological research project, initially presented in [6], involved examining fluctuating infant facial temperatures and detecting the exact location of temperature maxima. Without going into excessive detail regarding the project's medical aspects, it is theorised that a climax of increasing facial temperature closely correlates with other physiological episodes. Our task was to accurately determine the occurrence of these maxima in real-time.

Initially developed software implemented a simple hill climbing approach to determine the location (in time) of the temperature maxima. These times were correlated with the nearest occurrence of a particular physiological episode and the correlation value, or *p*-value, was 0.072. That is, the relationship was not significant.

The EC algorithm (width \times height) was also implemented to detect the temperature maxima. The same correlation calculation method was used and the *p*-value improved to 0.001 (highly statistically significant). These results appear to indicate that the EC algorithm is providing a more accurate estimate of the true temperature maxima.

V. CONCLUSIONS

This paper presents a novel local minima and maxima avoidance algorithm. The EC algorithm examines an extrema's

Technique	Error Rate
Simple Hill Climbing	6.9%
Thresholding	17.0%
Simulated Annealing	24.0%
Convolution and Hill Climbing	5.3%
EC	2.9%
Convolution and EC	3.4%
EC (width \times height)	2.3%

TABLE I

ERROR RATES USING VARIOUS METHODS OF OVERCOMING FALSE EXTREMA IN A SPECIFIC SIGNATURE VERIFICATION ENVIRONMENT. IF THERE ARE PARAMETERS INVOLVED IN THE OPERATION (SUCH AS CONVOLUTION WINDOW SIZE) THEN THE PARAMETERS PRODUCING THE LOWEST ERROR RATE WERE USED TO GENERATE THE RESULTS.

width or *consistency* more so than the actual height. Short-term peaks/valleys can be encountered in many situations (e.g., noisy environments/hardware inaccuracies) and can cause system accuracy problems. The EC algorithm is far less susceptible to these anomalies than existing techniques and tends to produce higher quality results.

The EC algorithm has shown that it can be effective in various practical iterative improvement environments. This paper discussed several specific large-scale applications, and performed a comparison between the EC algorithm and hill-climbing, convolution, thresholding and simulated annealing. The EC algorithm is a notable improvement over the other approaches in all of the explored applications.

Future work involves comparing the algorithm's effectiveness to other optimisation techniques such as genetic algorithms. Furthermore, the possibility of adapting the algorithm to multi-dimensional space will be explored. Should this be successful, it would be interesting to examine its utility in traversing error surfaces for more effective neural network training. We are also applying the EC algorithm to software bidding agents to enable correct measurements of maximum and minimum prices.

REFERENCES

- [1] Hirschman, I. and Widder, D. *The Convolution Transform*. Dover Publications, 2005.
- [2] McCabe, A. *Implementation and Analysis of a Handwritten Signature Verification Technique*. Honours Thesis, James Cook University, 1997.
- [3] McCabe, A. *Markov Modelling of Simple Directional Features for Effective and Efficient Handwriting Verification*. R. Mizoguchi and J. Slaney (Eds.): PRICAI 2000, LNAI 1886, pp 801(1-12), 2000.
- [4] McCabe, A. *Handwritten Signature Verification Using Complementary Statistical Models*. Ph.D. Thesis, James Cook University, 2004.
- [5] Plamondon, R. and Srihari, S. *On-line and Off-line Handwriting Recognition: A Comprehensive Survey*. TPAMI, Vol. 22, No. 1, pp 63-84, 2000.
- [6] Russell, M.J. and Vink, R. *Increased Facial Temperature as an Early Warning in Sudden Infant Death Syndrome*. Accepted for publication in Medical Hypotheses, 2001.
- [7] Russell, S. and Norvig, P. *Artificial Intelligence - A Modern Approach (2nd Edition)*. Prentice Hall, 2002.
- [8] Trevathan, J. and McCabe, A. *Remote Handwritten Signature Authentication*. In Proceedings of the 2nd International Conference on E-Business and Telecommunication Networks (ICETE '05), pp 335-339, 2005.

6.3 An Adaptive Shill Bidding Agent

This paper presents an adaptive shill bidding agent which when used over a series of auctions with substitutable items, can revise its strategy based on bidding behaviour in past auctions. The adaptive agent uses the EC algorithm (see Section 6.2) to determine the optimal price to aspire for. The adaptive agent then supplies the simple shill bidding agent (see Section 6.1) with the appropriate parameters which it uses to shill in an individual auction. The adaptive agent's ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. The performance of the EC algorithm is pitted against existing valley and peak detection algorithms. This paper illustrates the EC algorithm's usefulness in an additional application (i.e., auctions).

"An Adaptive Shill Bidding Agent" [16] is to appear in the International Conference on E-Business 2007 (ICE-B). This paper is available online via citeseer ⁴.

This paper is co-authored with Dr Alan McCabe of Fern Computer Services in Belfast, Ireland. Dr McCabe's research interests include Artificial Intelligence and Biometric Security.

⁴<http://citeseer.ist.psu.edu/>

An Adaptive Shill Bidding Agent

Jarrod Trevathan
School of Maths, Physics and IT
James Cook University
Email: jarrod.trevathan@jcu.edu.au

Alan McCabe
Fern Computer Services Ltd.
Belfast, Ireland
Email: alan@mymait.com

Wayne Read
School of Maths, Physics and IT
James Cook University
Email: wayne.read@jcu.edu.au

Abstract—This paper presents a software bidding agent that inserts fake bids on the seller’s behalf to inflate an auction’s price. This behaviour is referred to as shill bidding. Shill bidding is strictly prohibited by online auctioneers, as it defrauds unsuspecting buyers by forcing them to pay more for the item. The malicious bidding agent was constructed to aid in developing shill detection techniques. We have previously documented a simple shill bidding agent that incrementally increases the auction price until it reaches the desired profit target, or it becomes too risky to continue bidding. This paper presents an adaptive shill bidding agent which when used over a series of auctions with substitutable items, can revise its strategy based on bidding behaviour in past auctions. The adaptive agent applies a novel prediction technique referred to as the Extremum Consistency (EC) algorithm, to determine the optimal price to aspire for. The EC algorithm has successfully been used in handwritten signature verification for determining the maximum and minimum values in an input stream. The agent’s ability to inflate the price has been tested in a simulated marketplace and experimental results are presented.

I. INTRODUCTION

Agent based negotiation is now an integral part of online auctions and online share trading. An ever increasing number of transactions are being performed by automated bidding agents. However, limited attention has been paid to the security implications of using bidding agents, or the damage such agents can inflict when operating in an undesirable or fraudulent manner. Similar to a virus or worm, a *malicious bidding agent* can behave with an intent to do an auction harm. This might be in the form of inflating the auction’s price with fake bids (i.e., shilling), attacking the cryptographic protocols of a “secure” auction system, or launching a denial of service attack against the Auctioneer. In an extreme example, terrorists might unleash a malicious agent to trade in, and hinder the world’s stock exchanges, in an attempt to undermine the financial system. Alternately, they may attempt to obtain funds through fraudulent activities in auctions. Therefore the threat posed by malicious bidding agents to electronic commerce is very serious.

The Research Auction Server (RAS) ¹ at James Cook University, is an online server for conducting research into security issues regarding online auctions (see [9]). We are developing methods to detect shill bidding behaviour. Both real and simulated auctions are performed to test the effectiveness of the detection methods. To aid in testing, we developed a

software bidding agent which bids in a manner consistent with a shill. All bidding agents created thus far, operate in a controlled and near perfect environment. In contrast, human agents can be devious, which is reflected in real-world markets. As it is not permitted to use malicious agents on commercial online auctions, or in academic agent competitions (see [11]), we have created an agent interface for RAS. RAS allows the agents to be tested in a controlled (and legal) manner. This paper does not condone the use of these agents in any manner outside the scope of this research. By developing malicious bidding agents, we hope to better understand the characteristics of such agents, and how to protect against the damage they inflict.

We have previously documented a software bidding agent that follows a simple shill bidding strategy (see [10]). The agent incrementally increases the price during an auction, forcing legitimate bidders to submit higher bids in order to win an item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction (e.g., during slow bidding or near the auction’s end). Experimental results showed that the agent was able to inflate an auction’s average price by up to 25%, depending how risky the agent was prepared to be.

This paper presents an adaptive shill bidding agent. When used over a series of auctions with substitutable items, the adaptive agent is able to revise its strategy based on bidding behaviour in past auctions. The adaptive agent applies a novel prediction technique referred to as the Extremum Consistency (EC) algorithm, to determine the optimal price to aspire for. The EC algorithm has successfully been used in handwritten signature verification for determining an input stream’s maximum and minimum values in real-time. The agent’s ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. We show that the EC algorithm is superior to other valley and peak detection algorithms in an auctioning application.

This paper is organised as follows: Section II presents a simple shill bidding agent that shills in a single auction. Section III presents an adaptive shill bidding that uses the past history of auctions to predict and revise its strategy. Section IV evaluates the adaptive agent’s performance in terms of its ability to manipulate the auction in a manner that benefits the seller. Section V provides some concluding remarks.

¹<http://auction.math.jcu.edu.au>

II. A SIMPLE SHILL BIDDING AGENT

This section presents a simple shill bidding agent that uses bogus bids to inflate an auction's price. This section provides an insight into general shill behaviour. We briefly describe the agent's goals, how it interacts in the auction, and its strategic directives governing shilling (see [10] for a full description).

The main goal for shilling is to artificially inflate the price for the seller beyond what legitimate bidders would otherwise require to win the item. The seller's pay-off is the difference between the final price and the uninflated price. A shill's goal is to lose each auction. A shill is not constrained by a budget, but rather a profit margin. If the shill wins, the item is resold in a subsequent auction. However, there is a limit on how many times this can be done. For each auction a shill wins, the seller incurs auction listing fees and is required to invest more time. Continual wins erode the profit from shilling on the item.

The shill faces a dilemma for each bid they submit. Increasing a bid could marginally increase the seller's revenue. However, raising the price might also result in failure if it is not outbid before the auction terminates. The shill must decide whether to 'take the deal', or attempt to increase the pay-off. On the contrary, a bidder's goal is to win. A bidder has a finite budget and is after the lowest price possible. Increasing a bid for a legitimate bidder decreases the money saved, but increases the likelihood of winning. The following outlines typical shill behaviour and characteristics:

- A shill tends to bid exclusively in auctions only held by one (or a few) particular seller(s).
- A shill generally has a high bid frequency. An aggressive shill will continually outbid legitimate bids to inflate the final price. A shill typically will bid until the seller's expected pay-off for shilling has been reached. Or until the shill risks winning the auction (e.g., near the termination time or during slow bidding).
- A shill has few or no winnings for the auctions participated in.
- It is advantageous for a shill to bid within a small time period after a legitimate bid. Generally a shill wants to give legitimate bidders as much time as possible to submit a new bid before the closing time of the auction.
- A shill usually bids the minimum amount required to outbid a legitimate bidder. If the shill bids an amount that is much higher than the current highest bid, it is unlikely that a legitimate bidder will submit any more bids and the shill will win the auction.
- A shill's goal is to try and stimulate bidding. As a result, a shill will tend to bid more near the beginning of an auction. This means a shill can influence the entire auction process compared to a subset of it. Furthermore, bidding towards the end of an auction is risky as the shill could accidentally win.

A. The auction process and how the agent interacts in the proceedings

This paper restricts its attention to online English auctions resembling those commonly used online. However, many of

the principles can be extended to Continuous Double Auctions which are used in online share trading applications.

In order to participate in an auction, a bidder must *register*. They are provided with a unique bidder id, b_{id} , which they use to submit bids. During the *initialisation* stage, the Auctioneer sets up the auction and advertises it (i.e., item description, starting time, etc). An auction is given a unique number, a_{id} , for identification purposes. In the *bidding* stage, a bidder computes his/her bid and submits it to the Auctioneer. The agent can place a bid in auction a_{id} , for price p' , by invoking the **submit bid**(a_{id}, p') function.

The Auctioneer must *supply intermediate information* to the agent pertinent to the auction's current state. The agent can request a price quote for a particular auction by invoking the **obtain price quote**(a_{id}) function. This includes the start, end and current time for the auction, and the starting bid (if one exists). It is assumed that the agent has access to the entire bid history up to the current time in the auction. The history can be considered as an ordered set $\mathcal{H} = \{h_1, h_2, \dots, h_n\}$, $|\mathcal{H}| = n$, that contains price quote triples $h_i = (time, price, b_{id})$, where $1 \leq i \leq n$. The last element is the latest price quote for the auction (i.e., h_n is the current highest bid).

Finally, during the *winner determination* stage, the Auctioneer chooses the winner according to the auction rules (e.g., who has the highest bid, whether the reserve has been met, etc.).

B. Operation of the Simple Shill Bidding Agent

The agent's goal is to maximise the profit from shilling, while avoiding winning the auction. A shill that wins the auction is deemed to have failed. The shill agent bids according a strategy defined by a set of directives depending on the current auction state. Each directive plugs into the agent interface which dictates its bidding behaviour. The directives are as follows:

D_1 - Only bid the minimum amount required. As part of the price quote, the Auctioneer also provides the minimum amount required to out bid the highest bid. This is usually calculated as a percentage of the current high bid, or determined according to a scalable amount depending on the value of the current high bid. $D_1(p)$ is a function that takes the current price, p , and returns the minimum amount the shill should bid.

D_2 - Bid quickly after a rival bid. The agent must bid immediately in order to influence the other bidders for the maximum time.

D_3 - Don't bid too close to the end of an auction. If the shill bids too close to the end of an auction, it risks winning. To avoid this, the agent has a risk limit, θ . The agent is prohibited from bidding if the auction is more than $\theta\%$ complete. This is referred to as the *shill time limit*. Larger values of θ increase the risk that a shill might win an auction. $D_3(\theta)$ is a function that takes the risk limit θ , and returns true or false regarding whether the agent should continue to

bid.

D_4 - **Bid until the target price has been reached.** $D_4(\alpha)$ is a function that takes the current price p , the skill's target price α , and returns true or false regarding whether the agent should continue to bid. The agent will only bid when the current price p , is less than or equal to α .

D_5 - **Only bid when the current bidding volume is high.** The agent should preferably bid more towards the beginning of an auction and slow down towards the skill time limit, unless the bidding activity is high. That is, the volume of bidding must increase throughout the auction for the skill to maintain the same frequency of bidding. The agent uses the bid history \mathcal{H} , to analyse the current bid volume and decide whether to submit a bid. The agent observes the previous number of bids for a time interval. If the number of bids for the period is below a threshold, then the agent does not submit a bid. When no bids have been submitted for an auction, the skill agent will attempt to stimulate bidding by submitting the first bid. This is a common practice in auctions. It is intended that psychologically the presence of an initial bid raises the item's worth, as competitors see that it is in demand. $D_5(\mu)$ is a function that takes the risk limit μ , and returns true or false regarding whether the agent should continue to bid.

The following pseudocode illustrates the agent's behaviour:

```

skill agent(aid,  $\alpha$ ,  $\theta$ ,  $\mu$ ) {
  do {
    obtain price quote(aid)
    if ( $D_3(\theta)$  AND  $D_4(p, \alpha)$  AND  $D_5(\mu)$ )
      submit bid(aid,  $D_1(p)$ )
  } while ( $D_3(\theta)$  AND  $D_4(p, \alpha)$ )
}

```

The agent initially requests a price quote. If no bids have been submitted, then D_5 returns true and the agent submits a bid for the amount returned by D_1 . The agent then repeatedly requests price quotes to ensure that it is able to bid quickly if there is a rival bid (i.e., D_2). When a rival bid is submitted, the agent will bid only if the remaining directives D_3 , D_4 and D_5 are satisfied. The agent executes in this manner (i.e., requesting and evaluating price quotes), until either D_3 or D_4 becomes false.

III. AN ADAPTIVE SHILL BIDDING AGENT

In the case where there are multiple auctions for substitutable items (i.e., all items are the same), a skill agent can learn information that may help it be more successful over time. For example, if the final price for a series of auctions is constantly above the skill target price, then the agent can revise its target price upward. Alternately if the skill fails by not meeting its target, then it can revise the target price down. We refer to this as an *adaptive skill bidding agent*.

The adaptive agent is based on the simple skill agent and follows the same strategy. However, the adaptive agent

supplies the simple agent with differing risk values based on previous experience. This allows the simple skill agent to alter its strategy for each auction. This section describes the adaptive skill bidding agent's lifecycle, as well as the underlying prediction and revision techniques.

A. Approach

The adaptive skill agent operates in four phases: *preparation*, *planning*, *execution* and *revision*. Each of these stages are described in turn.

In the *preparation* phase, the agent is given a set of target auctions in which it will participate. The user provides the agent with the reserve price r , and a risk factor ϕ , $0 \leq \phi \leq 1$, which will dictate how much profit the skill can aspire to obtain. The agent is also supplied with bidding histories from similar past auctions. The *prediction method* uses the bidding histories to build a function, that given a bidding price, returns the probability that the price will win.

In the *planning* phase, the bidding agent sets the target price α , equal to the corresponding price with a probability indicated by the risk factor ϕ , according to the historical winning bid distribution. If $\alpha < r$, then the agent requests the user to lower r and/or increase ϕ . It is assumed that all auctions are held by one seller at the same auction house. If two auctions overlap (i.e., execute simultaneously), concurrently executing agents do not affect each other's operation.

In the *execution* phase, the adaptive agent executes the bidding plan by successively placing bids in each of the selected auctions (via the simple agent). The adaptive agent executes in a particular auction until the simple agent terminates (i.e., it reaches α or the time limit).

In the *revision* phase, the predictive bid function is updated with the auction's results depending on the agent's performance. Let ϕ' be the revised agent's risk factor, where $r \leq \phi' \leq \phi$. ϕ' is raised or lowered depending on the skill's success. The agent selects the next auction from the set of target auctions and re-enters the execution phase with α corresponding to ϕ' .

B. Prediction Methods

The adaptive agent constructs a probability function from the bidding histories of past auctions. In an English auction, the final price reflects the valuation of the second highest bidder. That is, the winner does not disclose the true highest amount they were willing to pay. Contrast this with First Price Sealed Bid (FPSB) and Vickrey auctions. In these auctions, bids are sealed so that a bidder does not know the value of anyone else's bid. The winner is the bidder with the highest bid. A FPSB auction requires the winner to pay an amount equal to the highest bid, whereas in a Vickrey auction, the winner pays an amount equal to the second highest bid. Constructing a probability function from a FPSB auction would yield an accurate depiction of the bidders' true valuation of an item. However, the Vickrey auction's probability function is the same as an English auction, in that only the winner's second highest valuation is possibly known.

Extrapolation techniques have been proposed to approximate the winner's true valuation from a set of past English auctions (see [3]). However, this information is not required by the shill agent. The shill's goal is to force the bidder into bidding their true valuation. Knowledge of the second highest price is satisfactory, as the shill can assume that by bidding somewhere within the range of the second highest price, that the buyer will be forced into inflating their bid nearer to his/her true valuation. It is too risky for the shill to try bid up to, or at the bidder's true valuation. Bidding at the second highest price is a much safer strategy, and should capture the majority of the desired profit from shilling.

[3] propose two methods that a bidding agent can use to construct a probability function. The first uses a histogram of the final auction prices, to be the function that maps a real number x , to the number of past auctions whose final price was exactly x . The final price of an auction a with no bids and zero reserve price, is then modeled as a random variable fp_a , whose probability distribution, written $P(fp_a = x)$, is equal to the histogram of final prices, scaled down so that its total mass is 1. The probability of winning an auction with a bid of z assuming no reserve price, is given by the cumulative version of this distribution, that is $P(fp_a \leq z) = \sum_{0 \leq x \leq z} P(fp_a = x)$, for an appropriate discretisation of the interval $[0, z]$. For example, if the sequence of observed final prices is $[20, 18, 23]$, the cumulative distribution at the beginning of an auction is:

$$P_a(z) = P(fp_a \leq z) = \begin{cases} 1 & \text{for } z \geq 23 \\ 0.66 & \text{for } 20 \leq z < 23 \\ 0.33 & \text{for } 18 \leq z < 20 \\ 0 & \text{for } z < 18 \end{cases}$$

In the case of an auction a with quote $q > 0$ (which is determined by the reserve price and the public bids) the probability of winning with a bid of z is:

$$P_a(z) = P(fp_a \leq z | fp_a \geq q) = \frac{P(fp_a \leq z \wedge fp_a \geq q)}{P(fp_a \geq q)} = \frac{\sum_{q \leq x \leq z} P(fp_a = x)}{\sum_{x \geq q} P(fp_a = x)}$$

In particular, $p_a(x) = 0$ if $z < q$.

[3] identify two drawbacks with the histogram method. First, the computation of the value of the cumulative distribution at a given point, depends on the size of the set of past auctions. Given that the shill agent heavily uses this function, this can create an overhead for large sets of past auctions. Second, the histogram method is inapplicable if the current quote of an auction is greater than the final price of all the past auctions, since the denominator of the above formula is then equal to zero. Intuitively, the histogram method is unable to extrapolate the probability of winning in an auction if the current quote has never been observed in the past.

The *normal method* addresses these two drawbacks, although it is not applicable in all cases. Assuming that the number of past auctions is large enough (more than 50), if the final prices of these auctions follow a normal distribution with mean μ and standard deviation σ , then the random variable fp_a can be given a normal distribution $N(\mu, \sigma)$. The

probability of winning with a bid z in an auction a with no bids and zero reserve price, is then given by the value at z of the corresponding cumulative normal distribution:

$$P_a(z) = P(fp_a \leq z) = \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\frac{z-\mu}{\sigma}} e^{-x^2/2} dx$$

If the current quote q of an auction a is greater than zero, the probability of winning this auction with a bid of z is:

$$P_a(z) = P(fp_a \leq z | fp_a \geq q) = \frac{P(fp_a \leq z \wedge fp_a \geq q)}{P(fp_a \geq q)} = \frac{\int_{\frac{q-\mu}{\sigma}}^{\frac{z-\mu}{\sigma}} e^{-x^2/2} dx}{\int_{\frac{q-\mu}{\sigma}}^{\infty} e^{-x^2/2} dx}$$

The complexity of these algorithms is only dependent on the required precision, not on the size of the dataset from which μ and σ are derived. Hence the normal method can scale up to large sets of past auctions. The normal method is able to compute a probability of winning an auction with a given bid, even if the value of the current quote in that auction is greater than all the final prices of past auctions. The domain of the normal distribution is the whole set of real numbers, unlike discrete distributions such as those derived from histograms. [3] performed an analysis of datasets from eBay and Yahoo² and showed that the final prices of a set of auctions for a given item are likely to follow a normal distribution. This is due to a given item having a more or less well-known value, around which most of the auctions should finish.

Figure 1 A illustrates how the adaptive agent performs on an example auction dataset using the aforementioned approach. The auction data is ordered according to time. A histogram is given showing the bid frequency. From this a cumulative probability distribution is derived. The adaptive agent is initially supplied with the risk parameter ϕ , $0 \leq \phi \leq 1$. The price corresponding to ϕ in the cumulative distribution is the maximum price the agent is willing to bid. α can be set to any price in the probability range up to ϕ . In this example, the agent's risk factor ϕ is 0.75, therefore the agent can set α up to \$6.85.

If the history of past auctions covers a large period of time, *data aging* must be taken into account. The normal method can be adapted to consider time-weighted averages and standard deviations of prices. In this way, recent observations are given more importance than older ones. Figure 1 B illustrates a time-weighted curve. In this example, the average time for a price observation is multiplied by its cumulative distribution and cumulatively summed. This value is then normalised against the maximum and minimum observations to give the time weighted cumulative distribution. The agent's risk factor ϕ remains the same 0.75, but α increases to \$8.48.

However, a time-weighted approach alone may not give the best results. Permanent price increases do occur over time due to inflation and other economic factors. As a result, past bids are no longer valid. Furthermore, temporary extreme price skews can also affect the process thus resulting in erroneous

²www.yahoo.com/auction




Fig. 1. Example illustrating how the adaptive agent basically plans and revises its strategy




Fig. 2. Example maximum situations encountered in an auction dataset

predictions. To address these problems, we propose that a valley/peak detection algorithm be used in determining the optimum target price.

Finding the true maximum/minimum price in real-time

The adaptive agent uses a novel approach for determining the correct price to set α in the presence of data aging and extreme price skews. This is referred to as the Extremum Consistency (EC) algorithm (see McCabe and Trevathan [6]). The EC Algorithm is used to find the true maxima or minima in a noisy input stream. The algorithm was initially developed for use in Handwritten Signature Verification. The problem arose when trying to detect a signature's turning points, and changes in other dynamic characteristics such as velocity and pressure in real-time. Noise is common in this environment due to hardware inaccuracies, which makes choosing the correct extrema difficult. The EC algorithm is able to avoid false extrema based on a tolerance parameter.

The adaptive agent uses the EC algorithm to determine the true maximum and minimum winning prices over a series of auctions. This allows the agent to set α according to the risk between these two extremes. As auction closing prices tend to oscillate around about a trend line, the agent needs the ability to distinguish meaningless price fluctuations from true extrema. The agent can update α between auctions when a new maximum or minimum is encountered (in real-time). For

example, when the price increases over time due to inflation.

Figure 2 presents several examples of maximum situations in encountered in auction datasets. Figure 2 A clearly has a single valid maximum. However B and C are most likely due to noise and should be ignored. D contains only a single valid maximum, where the dip between the two peaks is probably also due to noise (i.e., insignificant price fluctuations). The EC algorithm's goals in the auction application are:

- 1) Ignore meaningless price skews (*noise*).
- 2) Find the *true* maximum and minimum price in *real-time*.

The EC algorithm takes a different approach to other techniques such as hill climbing or convolution, in that it examines the width (or consistency) of an extremum rather than just its depth or height. This allows it to compensate for noise. In an auctioning application, noise to due to abnormal observations in the dataset such as an extreme price skew. Extreme price skews can occur for various economic factors. For example, supply of an item dramatically drops, whereas demand remains the same. This can occur when crops are destroyed by a storm, or war breaks out in the Middle East which restricts oil supply. Furthermore, an extreme price skew may occur when




Fig. 3. Example showing a local maximum and a true maximum

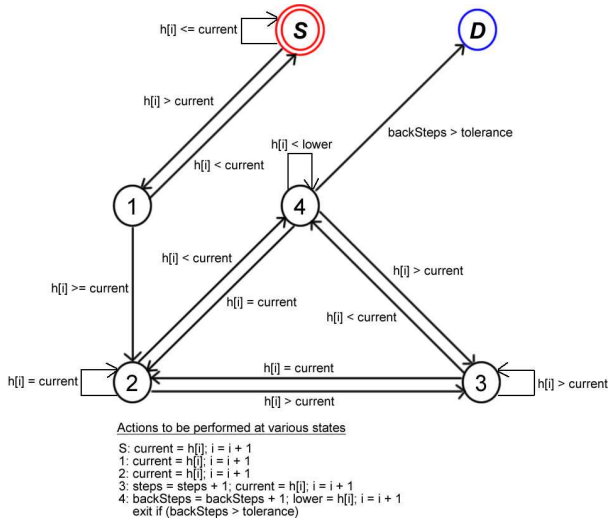


Fig. 4. A finite state machine expressing the EC algorithm for finding the “width” of, or number of “steps” in, the initial upslope of a peak. The movements between vertices (states) are defined by the comparison between points in the input stream and the comparisons are included on the edges in the diagram. Additionally there are actions to be performed when some vertices are reached - these are also included in the diagram. Note that in this table, $h[i]$ refers to the i^{th} element in the list of stream values h . Once the `backSteps` parameter exceeds the `tolerance`, the algorithm enters the dead state D and we have the width of the slope in the `steps` parameter. The width of the downslope is similarly calculable, with the inversion of various state transition conditions and the width of the valley itself is then the minimum of the downslope width and upslope width. Likewise, the width of peaks can be found with minimal modifications to the algorithm.

confidence in a particular item is momentarily low, such as an extortion threat to poison chocolate bars.

Figure 3 illustrates a dataset which contains two maximums, only one of which is valid. A local incorrect maximum is at $Time = 2$. In this case the peak itself is actually very short and using the maximum or mean slope size would make the peak appear erroneously large – the use of the minimum slope size is much more accurate. By placing a threshold on the width, this peak will ideally be ignored and that at $Time = 7$ would be taken as the more appropriate maximum. Figure 4 illustrates how the EC algorithm operates.

In shill application, the EC algorithm is used to determine which values should be included in the probability distribution. Initially, the known maximum and minimum are set equal to the first item in the dataset. Two instances of the EC algorithm are run concurrently. The first is referred to as EC_Max, and is tasked with searching for the latest maximum. Likewise EC_Min searches for the latest minimum. Once one of these algorithms terminates, this value is used as the current extremum and then the corresponding EC algorithm is executed. This process continues until the end of the dataset. For example, at the start of a dataset, if EC_Min terminates first, then EC_Max is prematurely terminated, its results discarded, and then run from the newly found minimum until a new max is found. Once the new max is found, EC_Min is executed starting from the new max. EC_Max and EC_Min alternate in this manner throughout the dataset.

Figure III-B illustrates the process of how the agent uses extrema to aid its predictions. In this example, the EC algorithm with a tolerance parameter of 3, reveals that there are three minima (\$1, \$5, \$5) and three maxima (\$8, \$10, \$8). The probability distribution only includes the values between a maximum/minimum pair. As the extremum are updated in real-time, this yields four different distributions for this example. We refer each maximum/minimum pair as an *epoch*. Once a new extremum has been determined, the agent uses the most recent epoch and ignores all other past values. Each new value encountered in the yet undetermined epoch is also incorporated into current epoch’s probability distribution. On discovery of the next extremum, the epoch currently in use is discarded and the newly discovered epoch used.

The tolerance parameter is altered according to the outlook. A small tolerance is used for short term predictions and high tolerance for long term predictions. In the ongoing example, increasing the tolerance results in the minimum and maximum being determined as \$1 and \$10 respectively.

C. Revision Strategies

In the revision phase, the agent assesses its performance and revises its strategy. If the agent is successful in that it did not win the auction and achieved its target, the value of the winning bid h_n , is incorporated into the current epoch. ϕ' is revised upwards by a factor ϵ , where $\phi' + \epsilon \leq \phi$.

If the agent has failed by winning the auction, then ϕ' is revised downward by a factor ϵ , where $r \leq \phi' - \epsilon$. When the skill fails in this manner, the second highest bid (i.e., h_{n-1}), is added to the current epoch and the skill’s winning price is discarded. If the agent did not win, and failed to achieve its target, ϕ' remains the same. The value of the winning bid h_n , is still incorporated into the current epoch.

Determining the agent’s risk profile

A riskier shill can achieve more profit, but at the increased likelihood of winning an auction. Directives D_3 and D_4 determine whether the agent should bid based on the current auction time and bid volume.

The agent can also revise other risk factors in response to historical bidding patterns. θ can be increased if bidding in previous auctions has tended to occur closer to the end of an auction. Likewise, μ can be increased if the bid frequency in previous auctions is low.

IV. PERFORMANCE

This section describes the performance of the adaptive shill bidding agent. Claims regarding the simple shill agent can be found in [10]. The adaptive agent is implemented on RAS and has been tested with other types of bidding agents in a simulated auction market. The adaptive shill agent is assessed on its ability to inflate an auction’s final price. A shill bidding agent is considered successful if the final price equals or exceeds the shill’s target price. The shill agent is considered unsuccessful if it won the auction, or failed to reach the target price.




Fig. 5. Example illustrating how the EC algorithm influences the agent’s probability function. Epochs are created between max/min pairs in the time series chart (top). Prices further in past are discarded. The agent uses the price corresponding to ϕ for the current epoch in use.

A. Elements of the experimental setup

Seed data a dataset was obtained from a series of simulated auctions. This was used as a “seed” to initialise the skill agent (i.e., provide it with an initial epoch).

Zero Intelligence (ZI) Bidder This bidding agent is designed to simulate an ordinary bidder in an auction. It is assigned a random amount which it tries to submit as a proxy bid at a random time throughout the auction. A ZI bidder’s limit price is generated randomly based on the a trend factor. Specifically, a predetermined trend was programmed which randomly assigns the agent’s limit price between a given maximum and minimum. The extremum were increased and decreased over time to give the appearance of natural price trends (compared to completely random prices).

Shill Bidder This is an implementation of the adaptive skill bidding agent proposed in this paper. The shill agent has a reserve price r , and a profit risk factor ϕ . The agent’s target price α , is determined according to the current epoch, and its strategy is adjusted using the aforementioned revision technique.

Auction A software simulated auction. The auction has a start and end time. The bidding agents can submit bids during this time, where the auction outcome depends on the auctions of the agents. All auctions are English auctions with proxy bidding.

B. Claims, experiments, and results

Claim 1 *The adaptive skill agent achieved a higher success rate than the simple skill agent.* To validate this claim, we pitted the adaptive agent against the simple agent using the same dataset and risk parameters. On average (for the specific dataset), the adaptive agent achieved an 82% success rate, whereas the simple agent achieved a 43% success rate. The amount of profit acquired by the adaptive agent was 36% greater than the simple agent. This seems to indicate that it is worthwhile employing the adaptive agent.

Claim 2 *A riskier agent achieved a higher average final price than an agent following a safer strategy.* The agent was run on the same dataset with increasing values of ϕ . The riskier agent was able to achieve a 15% increase in the average final price compared to a risk-averse agent. However, this came at the expense of an increased failure rate due to winning.

Claim 3 *EC produced better results than simple gradient descent, thresholding and convolution.* To improve the agent’s profit, two alternate algorithms were implemented: simple removal of small valleys or peaks (called “thresholding”) and basic convolution of the data prior to gradient-descent/hill-climbing. Thresholding involved determining the distance between the peak’s location and the preceding valley’s location (that is, the depth of a valley or height of a peak). If this distance was below a specified threshold then that peak was ignored and the traversal continued in the same direction.

Examples of situations where thresholding was successful can be seen in Figure 2 B and C. However, the agent’s overall performance (in terms of profit acquired) was quite

Technique	Avg Final Price
Simple Hill Climbing	\$ 4.80
Thresholding	\$ 5.11
Convolution and Hill Climbing	\$ 6.20
EC	\$ 6.90

TABLE I
PROFIT ACQUIRED BY SHILLING.

poor. The reason it seems, is that a peak's *height* alone, while obviously containing useful information, is not the best validity indicator (at least in this environment), but rather the consistency/duration is more important.

Convolution is a method for "smoothing out" or "averaging" one-off "bumps" or random noise while attempting to preserve those extrema which are truly indicative of the price trend direction. The basic idea behind convolution is that a window of some finite length is scanned across the stream of values [1]. The output price is the weighted sum of the input prices within the window where the weights can be adjusted to perform various filtering tasks - when smoothing is performed the weights are generally all equal. After the input stream was convoluted the hill-climbing/gradient-descent approach was used to obtain the extrema.

Table I summarises the error rates of the implemented approaches used by the adaptive shill agent. The EC algorithm clearly achieves more significant results.

The EC algorithm's main advantage over convolution is execution speed. In a real-time application such a continuous double auction, execution speed can become a serious issue. In order to perform convolution an entire extra layer of computation is required, as convolution of the raw data must be done prior to obtaining the extrema, whereas with the EC algorithm the checking is done at the same time as the search for extrema. Additionally, convolution can become quite expensive using a large window or with a large raw data size.

Empirical experimentation with the adaptive shill bidding agent has found that convolution causes an average slowdown of 20-25% (depending on system parameters). The number of extra calculations required in the EC algorithm compared to naive hill-climbing is almost negligible with the slowdown of the shill agent experimentally found to be less than 3%.

Claim 4 *The adaptive agent's prediction method can be reasonably applied to share market prediction.* The prediction method was run on a data obtained from the Australian Stock Exchange. As previously mentioned, shares are traded in an auction type referred to as a Continuous Double Auction (CDA). A CDA has many buyers and sellers continually trading a commodity. Rather than shilling, the prediction method was used to determine maximum and minimum cycles in the closing prices for several listed companies. While the prediction method can not exactly determine when a new extremum will be encountered, it definitely can tell when an extremum has occurred and which direction the current price is moving. In this case, we programmed the agent to purchase shares immediately after a minimum had been determined,

and sell these once the corresponding maximum had been found. We also reversed this process. Depending on the its risk, for most datasets the agent was able to make a profit (i.e., finish with a greater value than initially provided). It is interesting to note that the latter approach (i.e., purchasing shares after the maximum and selling upon determining the minimum) produced better results.

V. CONCLUSIONS

This paper presents an adaptive shill bidding agent. When used over a series of auctions with substitutable items, the adaptive agent is able to revise its strategy based on bidding behaviour in past auctions. The adaptive agent applies a novel prediction technique referred to as the Extremum Consistency (EC) algorithm, to determine the optimal price to aspire for. The EC algorithm has successfully been used in handwritten signature verification for determining an input stream's maximum and minimum values in real-time. The agent's ability to inflate the price has been tested in a simulated marketplace and experimental results are presented. We show the superiority of the EC algorithm over other valley and peak detection algorithms.

In future work it would be useful to investigate agents that attack security protocols or launch denial of service attacks against the Auctioneer. Furthermore, we plan to devise a shill bidding agent that profiles other bidders based on their bidder id. The agent alters its strategy in response to the bidder's anticipated strategy based on previous observations. Finally, similar problems to shilling occur in CDAs such as "ramping the market". It would be intuitive to investigate agents that conduct this kind of fraud.

REFERENCES

- [1] Hirschman, I. and Widder, D. *The Convolution Transform*. Dover Publications, 2005.
- [2] Cliff, D., Minimal-intelligence agents for bargaining behaviours in market-based environments, *Hewlett Packard Labs, Technical Report HPL-97-91*, 1997.
- [3] Dumas, M., Aldred, L. and Governatori, G., A probabilistic approach to automated bidding in alternative auctions, In *Proceedings of the eleventh international conference on World Wide Web*, 99-108. ACM Press, 2002.
- [4] Gjerstad, S. and Dickhaut, J., Price formation in double auctions, *Games and Economic Behavior*, 22, 1-29, 1998.
- [5] Gode, D. and Sunder, S., Allocative efficiency of markets with zero intelligence traders: Market as a partial substitute for individual rationality, *Journal of Political Economy*, 101, 119-137, 1993.
- [6] McCabe, A. and Trevathan, J., A new approach to avoiding the local extrema trap, In *Proceedings of the 13th International Computational Techniques and Applications Conference*. 2006.
- [7] Rust, J., Miller, J. and Palmer, R., Behaviour of trading automata in a computerized double auction market, In *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1992.
- [8] Schwartz, J. and Dobrzynski, J., 3 men are charged with fraud in 1 100 art auctions on eBay, *The New York Times*, 2002.
- [9] Trevathan, J. and Read, W., RAS: a system for supporting research in online auctions, *ACM Crossroads*, 12.4, 23-30, 2006.
- [10] Trevathan, J. and Read, W., A simple shill bidding agent, In *Proceedings of the 4th International Conference on Information Technology - New Generations*, (to appear), 2007. Available at <http://auction.maths.jcu.edu.au/research/agent.pdf>
- [11] Wellman, M. and Wurman, P., The 2001 trading agent competition, *Electronic Markets*, 13.1, 4-12, 2003.

Chapter 7

Privacy and Security in Continuous Double Auctions

This chapter addresses privacy and security issues inherent in conducting online CDAs. Online share trading is an off-shoot of online auctions that allows an individual to buy and sell securities using an online broker. Little research has previously been conducted on security in this auctioning application. Two papers are presented. The first describes the basic security requirements for online share trading (see Section 7.1). The second presents a scheme for conducting anonymous and secure CDAs (see Section 7.2).

7.1 Privacy and Security Concerns for Online Share Trading

This paper provides an overview of privacy and security for online share trading in the context of auctions. The additional parties involved are described, and how their individual requirements affect auction privacy and security. This paper gives a description of share market related crimes and their implications for online trading. It also investigates existing security mechanisms and cryptographic solutions to these problems, and presents an auction scheme for anonymous and secure share transactions. This paper is designed to provide a general introduction to the paper of Section 7.2. It shows how the proposed CDA scheme fits into the broader application of online share trading.

“Privacy and Security Concerns for Online Share Trading” [17] is currently in submission with IEEE Security and Privacy. The magazine is peer-reviewed and published quarterly.

Privacy and Security Concerns for Online Share Trading

Jarrod Trevathan
School of Mathematics, Physics and IT
James Cook University, North Queensland, Australia
Email: jarrod.trevathan@jcu.edu.au

Abstract—Online share trading has made investing in share markets more accessible to novice traders. Investors can trade from the comfort and privacy of their own homes. However, there are many inherent risks in online share trading. Share markets are becoming the target of new elaborate scams, terrorist attacks and corporate crimes. This article provides an overview of privacy and security for online share trading in the context of auctions. We discuss the additional parties involved, and how their individual requirements affect auction security and privacy. We give a description of share market related crimes and their implications for online trading. We investigate existing security mechanisms and cryptographic solutions to these problems. This article presents an auction scheme for anonymous and secure share transactions.

I. INTRODUCTION

Brokers (such as Commsec ¹) now offer a multitude of online services to small traders for accessing the world's financial markets. Online traders perform an ever-increasing number of share market transactions. Share markets such as the New York Stock Exchange ² and the Australian Stock Exchange ³ use an auctioning mechanism referred to as a Continuous Double Auction (CDA). This is an auction that has many buyers and sellers continuously trading a commodity. In an online share trading system, traders are the bidders and the Share Market is the Auctioneer. However, unlike a regular auction, bidders must submit their bids to the Auctioneer via a Broker. The Broker earns a commission for its services.

Despite online share trading's popularity, there are security and anonymity concerns regarding the integrity of the participants, and the auctioning process. Share markets are tightly regulated. However, many of the problems inherent in online share trading, are similar to the problems encountered by more conventional online auctions such as eBay ⁴. For example, a trader might repudiate having made a bid, or refuse to pay for/deliver shares. Furthermore, the Broker/Auctioneer could be corrupt and initiate unauthorised trades. In addition, much information is revealed about an individual, including his/her personal information (identity, address, etc.), and trading history. Such information could potentially be used against him/her, or for marketing purposes.

¹<http://www.comsec.com.au>

²<http://www.nyse.au>

³<http://www.asx.com.au>

⁴<http://www.ebay.com>




Fig. 1. The online share trading process

This article provides an overview of online share trading in the context of auctions. We discuss the additional parties involved, and how their individual requirements affect auction privacy and security. We provide a description of share market related crimes and their implications for online trading. We also investigate existing security mechanisms and cryptographic solutions to these problems. Finally, we present a scheme for anonymous and secure share transactions.

II. ONLINE SHARE TRADING

Figure 1 illustrates the major parties involved in the share trading process. In order to trade shares, a Client must first obtain an account with a Broker. This is usually done offline via a form. Once registered, the Broker issues the Client with a Client Number. This number is unique to the Client (i.e., no two clients have the same number). The Client enters his/her bid via computer. This is communicated to the Broker's computer and is then entered into the Share Market.

This is analogous to online auctions where the Trader is a bidder, and the Share Market is the Auctioneer. However, unlike a regular auction, a bidder submits a bid to the Auctioneer via an intermediary (i.e., the Broker). Furthermore, the Share Market is a more complicated entity than a regular Auctioneer. The Share Market consists of numerous parties with differing interests in the auction proceedings.

Figure 1 also shows the additional parties involved in trading shares. A Company issues shares to the Share Market in order to raise capital. The Company releases statements to the Media, such as profit reports or market sensitive announcements. The Public observes the Share Market for potential investment opportunities, or for entertainment. It is also the Broker's

Buy

Order Details

Account No: 413678

ASX Code: URL

OR

Security Name: UNIVERSAL RESOU

At Market **OR** At Limit: \$ 0.160 (per share)

Quantity: 22000

Your Contact Details

First Name: Jarrod

Last Name: Trevathan

Telephone Number: 47731343

Order Estimate

Security: UNIVERSAL RESOURCES FPO

Quantity: 22000

Indicative Price: \$0.160

SubTotal: \$3520.000

Brokerage & Costs Incl GST: \$19.950

Stamp Duty: \$0.000

Indicative Total: \$3539.950

Indicative Price as at 20:48 Sydney Time, Friday, 7 April 2006

Basis of Quotation:

Fig. 2. An example buy bid used in online share trading

role to provide clients with information regarding the market, and/or advice concerning when to buy or sell.

At the end of a completed transaction, the Broker debits or credits a Client's bank account. The Share Registry is a trusted authority tasked with keeping records on share transactions and ownership. The Government is able to use the Share Registry to ensure that bidders have paid taxes on their transactions, and to check compliance with disclosure and ownership laws. Furthermore, companies require some record of who their shareholders are. This is to facilitate dividend payment, allow voting at shareholder meetings, and dissemination of market information.

Brokers allow clients to enter new bids and to modify/cancel existing bids. Figure 2 shows a typical web form a Client would use to submit a buy bid. (Note that a sell bid is just the opposite.) First a Client provides his/her Client Number, so that the Broker can link the order they submit to them. Next, the Client enters either the company code, or the full name of the company they want to purchase shares in. The company code is unique to a company, used for identification purposes.

There are two options for entering the price at which the Client wishes to purchase shares. With an "at limit" order, the Client enters the exact price they are willing to buy the shares at. When the Broker acts on the order, they will purchase

the shares at a price no higher than the limit the Client has specified. Alternately, an "at market" order instructs the Broker to purchase the shares from the next available seller, regardless of the price. Next, the Client enters the quantity of shares they want to buy. The Client is also required to enter some personal information about themselves including first name, last name and phone number.

The right side of Figure 2 gives an estimation of the Client's buy order. For this example, the Client has specified that they want to buy shares in Universal Resources (URL). The Client desires 22,000 units at a limit of \$0.16, thus the basic trade's total value is $22,000 \times \$0.16 = \$3,520.00$. The Broker also adds commission of \$19.95. This brings the total price that the Client must pay to \$3,539.95.

Figure 3 shows a market depth indicator, which lists all the buy and sell bids grouped according to price. The quantity is the aggregate of all bids at the particular price level. Every time a Client submits a new order, this information is immediately reflected via the market depth. Market depth information is available to all traders. When a trade is executed, the Broker notifies the Client via e-mail.

III. PRIVACY AND SECURITY ISSUES INHERENT IN TRADING SHARES ONLINE

Auction security issues have been fully discussed in [2], [3], [4], [7], [8]. Shares are traded in an auction type referred to as a Continuous Double Auction (CDA). Therefore, online share trading exhibits similar problems to conventional online auctions. However, a CDA is much more complicated, as there are many buyers and sellers continuously trading a commodity. This requires a more sophisticated clearing process (i.e., the manner which buy and sell bids are matched). Furthermore, the additional parties involved in the share trading process make it difficult to clearly define privacy and security requirements.

A. Privacy Concerns

In existing 'secure' auction schemes, privacy is regarded as a fundamental design goal. This is due to profiling, where the Auctioneer or seller uses a Client's information to force them into paying a higher price. For example, if a bidder's true valuation is discovered, then the seller may set the auction's reserve price at this amount. Alternately the seller may engage in shill bidding to artificially inflate the auction's price to the bidder's valuation thus ensuring maximum profit.

Online Brokers have largely neglected anonymity issues. The current approach lets the Broker and share market learn all information regarding the Client and his/her order. This is undesirable as personal information can be abused, or used against an individual.

If a Client has a high profile, they might want to conceal the fact that they are bidding. This might occur in the case where the Client feels that their privacy would be compromised in some manner (e.g., reveal that they have a fetish for an item), change other bidders' perceptions of the item (e.g., stimulate unwanted bidding), or influence the seller to raise the commodity's price.

◀ UNIVERSAL RESOURCES FPO ▶

Stock	Bid \$	Offer \$	Last \$	Change *	Open \$	High \$	Low \$	Volume	News
URL	0.155	0.160	0.160	0.000	0.170	0.170	0.150	1,436,631	

[Buy](#) | [Sell](#) | [Add to Watchlist](#) | [Research](#) | [Chart](#) | [Print](#) | [Help](#)

Market Depth**

BUY				SELL			
Number	Quantity	Price	#	Price	Quantity	Number	
1	58,000	0.155	1	0.160	92,852	5	
11	478,315	0.150	2	0.165	320,000	6	
2	150,000	0.145	3	0.170	185,727	3	
8	351,857	0.140	4	0.175	135,000	2	
2	120,000	0.135	5	0.180	67,000	2	
2	138,018	0.130	6	0.185	125,000	2	
2	49,800	0.125	7	0.190	615,523	6	
9	160,000	0.120	8	0.195	42,000	2	
5	600,900	0.115	9	0.200	438,300	5	
6	773,000	0.110	10	0.210	50,000	1	

Fig. 3. An example market depth indicator which lists all the buy and sell bids for a CDA

It is common for companies to sell client information to marketing agencies. Such information can include names, addresses, account and portfolio balances, and historical trading data. This often results in an individual being profiled and targeted with junk mail. For example, a Client may continually be solicited by offers for credit cards, investment trusts, or share market tip-sheets.

Furthermore, shareholders could be harassed during disputes involving the company. For example, during strikes, or if the company is engaging in unethical behaviour. Unethical behaviour may include pollution or selling products that cause harm (e.g., cigarettes, weapons, products containing asbestos, etc.). Furthermore, in the event of a hostile takeover, shareholders of the company being taken over may be harassed by the takeover advocates. It should be up to an individual whether or not they want to be associated with his/her share holdings in these situations.

However, the need for privacy must be balanced against the propensity for clients to cheat. Identity escrow (or recovery) is required in the situation where the Client has acted illegally. In this situation, it is often necessary for regulatory authorities to trace past transactions. Furthermore, there are purposes related to Government and auditing procedures that may require an individual's identity to be recovered.

B. Security Concerns

The most pressing security problem in online share trading is that there is no means for a Client to verify whether his/her order has been submitted to the Share Market. A dubious Broker might alter a Client's order (thereby defrauding the Client), or block it in the case where s/he disagrees with the Client's judgement, or has a vested interest in some aspect of the trade. As the Broker gains a commission proportional to the quantity traded, they might buy or sell a larger portion than warranted.

A further concern is that there is no evidence that the Share Market has included the order in the auction. Nor is there any verification of how the orders are being matched. Bidders must rely on market depth indicators. However, there is presently no means of recourse if the bid does not appear, nor if it was matched in a manner inconsistent with the clearing rules.

The Broker and Share Market must also be able to verify that the Client submitted the order. Otherwise, the Client may repudiate having made the bid. Furthermore, lack of authentication makes it possible to forge an order and hence frame innocent clients. For example, the Broker or Auctioneer may forge an order on a Client's behalf, or an outsider that wants to cause financial hardship to a Client.

Online share broking schemes use various solutions to enforce payment including credit, client accounts and fines. Credit allows clients to submit buy bids up to a set credit limit. If the bid is cleared, the Broker requests payment from the Client's Bank. Another solution requires clients to hold accounts with the Broker, and only allow bids up to the value of the account balance. When a Client submits a buy bid and this bid is cleared by the auction, the bid value is

withdrawn from the Client's account. Alternately, if a sell bid is cleared, the value of the bid is deposited in the Client's account. Finally, Brokers can impose fines for defaulting on payment, and can sell winnings in an attempt to recuperate unpaid money.

However, none of these solutions are perfect. Credit limits provide the Client with less restrictions, but requires the Broker to place some confidence in the Client's ability to pay. A Broker held account prevents risk on the Broker's behalf, but can be restrictive to the Client who might want to hold the account funds elsewhere while not in use. For example, if a Broker offers a lower rate of interest compared to another financial institution, the Client might want to move the funds to the better paying account when the funds are not being used. Fines and selling off goods won is a last ditch effort only taken to either deter or recover lost money after a Client has defaulted on payment.

An additional concern is that a Broker might not debit/credit the Client's bank for the correct amount. While this may appear to be extreme, and the reader may be thinking, "this wouldn't happen in real life", consider what would happen if it did. What means of recourse does a Client have? Essentially it would be the Client's word against the Broker's word, and it would require auditing the trade history and litigation to resolve. Once again, this comes down to a Client's ability to authenticate his/her bids and verification of the auction proceedings.

C. Share Market Crimes

The share trading process has many unique crimes compared to conventional online auctions. *Insider trading* generally refers to buying or selling a security (shares), based on the knowledge of non-public information. The insider has a fiduciary duty or relationship of trust and confidence regarding the security. For example, a company CEO sells all his/her shares prior to, and with knowledge that the company is about to go bankrupt. In this situation the CEO has acted illegally, as the Share Market does not yet know such information. A recent high profile individual to be convicted in the U.S. of insider trading is Martha Stewart ⁵.

Insider trading violations may also include "tipping" such information, securities trading by the person "tipped", and securities trading by those who misappropriate such information. Rene Rivkin ⁶ (in Australia) was convicted of insider trading when he acted on a tip provided by a QANTAS executive. Rivkin allegedly sold his QANTAS share holdings in anticipation of bad market related news, thus avoiding any loss when the news later became public.

Other share market crimes involve *misinformation*, where a company influences its share price by announcing incorrect statements regarding its financial standing. Furthermore, Brokers or financial advisers might be in possession of shares, which they sell to their clients for a profit using misinformation. "Ramping the market" refers to actions designed to

⁵<http://www.sec.gov/news/press/2003-69.htm>

⁶<http://www.wsws.org/articles/2003/jun2003/rivk-j07.shtml>

TABLE I
REGULATORY AUTHORITIES AROUND THE WORLD

Country	Authority Name
Australia	Australian Securities and Investments Commission
Canada	Ontario Securities Commission
China	China Securities Regulatory Commission
Israel	Israel Securities Authority
France	Commission Des Operations De Bourse
Germany	Bundesaufsichtsamt Fur Den Wertpapierhandel
Japan	Securities Bureau of the Ministry of Finance
Russia	Federal Commission for the Securities Market of the Russian Federation
United Kingdom	Financial Services Authorities
USA	U.S. Securities and Exchange Commission

artificially raise the market price of listed securities and to give the impression of voluminous trading, in order to make a quick profit. Once the bidding slows, the price falls back to its original, thus depriving those who purchased shares during the hype.

D. Existing Mechanisms for Ensuring Privacy and Security

Existing online trading sites only require clients to log-in using a Client Number and password. Once logged in, application level encryption is used to secure the session. The log-in usually times out after a period of inactivity, which is useful in the case where a Client leaves his/her computer for a long time, or forgets to log-off. Privacy is enforced by a privacy agreement among the Broker and Client. The Government provides policing in the form of regulation.

A Regulatory Authority is a law-enforcement entity that governs share market behaviour. Table I lists some of the World's major authorities. These authorities have the power to prosecute, incarcerate, fine or ban individuals guilty of share market crime. Such authorities were originally created to regulate conventional offline trading. Online share trading's similarity to online auctions means that authorities must be capable of preventing auction-related fraud. However, this seems doubtful if commercial online auctioneers are unable to solve the problem.

Further concerns also regarding trust cast doubt over a Regulatory Authority's effectiveness. In reality, a Regulatory Authority is just another individual that is involved in the auction process. They are no more trustworthy than anyone else. Some individuals have even felt that authorities have had unfair bias or vendettas against them. Furthermore, such authorities are too expensive to maintain and to seek recourse through. The authority must be funded or raise money to continue running. This ultimately makes the authority biased towards the funding source. Additionally, they are only effective to police large-scale fraud, whereas, the process is too complicated and expensive for small fraud instances. The question could be asked, "Are they doing enough to protect the privacy and rights of individuals"?

IV. CRYPTOGRAPHIC SCHEMES FOR CONDUCTING ANONYMOUS AND SECURE CDAS

Cryptography can be used to provide bid authentication and privacy in auctions. Most existing auction security proposals

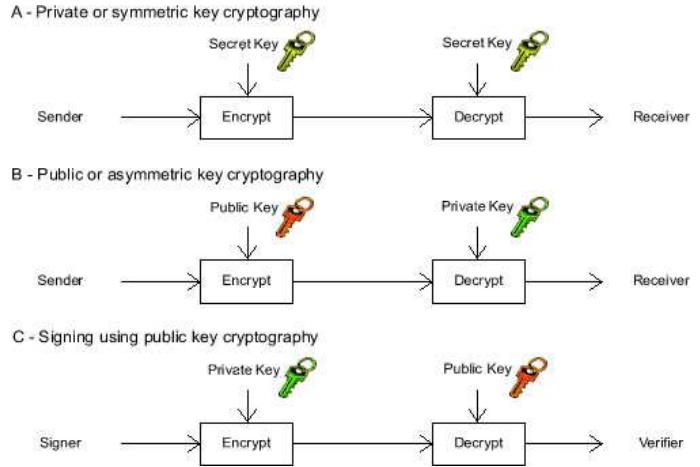


Fig. 4. Public/Private Key Encryption and Digital Signatures

have focused on sealed bid auctions (e.g., see [2], [8]). Wang and Leung [9] proposed a scheme for private and secure CDAs. However, this scheme has several security concerns and is not specific to share markets. Trevathan *et al* [5] propose a CDA scheme, which uses a cryptographic mechanism referred to as a *group signature*. A group signature alleviates many of the security and anonymity problems described in the previous section. To understand how a group signature works, we first provide some background on more general cryptography.

Cryptography allows two parties to communicate over an insecure medium (such as a network). This is achieved by encrypting a message using an encryption algorithm and a key. The message is decrypted at the receiving end using a decryption algorithm and the key. As long as an eavesdropper doesn't know the key, he/she can't read the message.

There are two main types of cryptography: private and public. *Private key* cryptography uses a single key to encrypt and decrypt a message. When encrypting/decrypting a message, both sender and receiver must possess the secret key. This process is shown in Figure 4 A.

Alternately, *public key* cryptography uses two keys, one for encryption and the other for decryption. The public key is made available to everyone and is used to encrypt messages. The private key is only known by the receiver and is used to decrypt messages. In general, it is hard to learn any




Fig. 5. Group Signature Scheme

information about the private key using the public key. This process is shown in Figure 4 B.

A digital signature allows the sender to sign a message such that the receiver can authenticate that it originated from the sender. This provides proof of the signer’s identity and prevents repudiation of the signed message. Public key cryptography can be used to sign messages (Figure 4 C). The sender signs the message using his/her private key. The receiver can then verify the signature using the sender’s public key. As only the sender knows his/her private key, the message could have only originated from him/her.

A group signature scheme allows group members to sign messages on the group’s behalf. Anyone can verify that the signature is authentic using the group’s public key. However, it is not possible to determine which group member signed the message. To join the group, a user must register with a Group Manager (GM) (see Figure 5). The GM is responsible for issuing group members with a private key and maintaining the public key. In the case of a dispute (e.g., a member denies having signed a message) the GM can trace the signer’s identity.

Group signature schemes naturally lend themselves to auctions. In this case a bidder belongs to a group of bidders. A bidder can sign bids on the group’s behalf in such a manner that s/he remain anonymous. The Auctioneer can verify the signature on a bid using the group’s public key. The group signature prevents bids from being forged, and allows the bidder’s identity to be revealed if a bid is repudiated. However, in terms of an auction, the Auctioneer is not the appropriate choice for the GM. The GM is generally trusted and therefore the Auctioneer cannot perform this role, as they could cheat by revealing a bidder’s identity without due cause.

The Trevathan *et al* [5] CDA employs a group signature scheme to allow for anonymous bid submission. The scheme introduces a Registrar that is responsible for registering bidders. Figure 6 shows the modified group signature scheme. A bidder registers with both the Auctioneer and Registrar in a manner such that his/her identity is split between the two parties. Neither the Auctioneer nor Registrar individually knows the bidder’s identity. A bidder remains anonymous as long as both parties don’t collude.

The CDA protocol is presented in Figure 7. The arrows represent the messages sent between parties. Dashed circles




Fig. 6. An auction scheme based on a Group Signature Scheme

surrounding messages indicate that the enclosed messages are sent during a particular stage of the auction protocol. The basic protocol for the CDA scheme is as follows:

- 1) *Registration*: A bidder first registers with the Auctioneer supplying his/her ID. The Auctioneer issues the bidder with a signed token. The bidder presents this to the Registrar. The Registrar verifies the signed token and issues the bidder with a group certificate.
- 2) *Bidding*: A registered bidder signs a bid using his group certificate. The bid is then submitted to the Auctioneer. The Auctioneer checks the bid to ensure that it is correctly submitted by a legitimate bidder. Signed bids are posted on a public bulletin board. This can be considered as a web page that every one has read permission, but only the Auctioneer can write to. This allows participants to observe and verify the auction proceedings.
- 3) *Winner Determination*: The Auctioneer determines the outcome of the auction according to the auction’s rules. A bidder can verify they have won by reproducing the signature on the winning bid.
- 4) *Tracing and Revocation*: The Auctioneer and the Registrar combine their information to determine a bidder’s identity and if required, to revoke them from the auction indefinitely.

By incorporating a group signature scheme, a CDA can achieve the security characteristics in Table II.

V. SECURE AND ANONYMOUS ONLINE SHARE TRADING

The Trevathan *et al* [5] scheme is specific to CDAs and does not address concerns involving the Broker. In this section we define the Broker’s role in relation to a Client’s privacy and




Fig. 7. The basic CDA protocol

TABLE II
SECURITY AND PRIVACY GOALS FOR A CDA

Unforgeability	Only registered bidders can issue a bid.
Anonymity	The identity of a bidder remains anonymous (i.e., a bidder cannot be associated with the bid s/he has submitted).
Unlinkability	No party can form a profile about a bidder's bidding strategy based on past bids.
Traceability	In case of dispute, the Registrar and Auctioneer, cooperatively can identify the bidder.
Exculpability	The Registrar and Auctioneer cannot issue a bid on behalf of any other bidder.
Coalition-resistance	A colluding subset of registered bidders cannot issue a bid that cannot be linked back to one of the colluding bidders.
Verifiability	All auction participants can be verified as having followed the protocol correctly.
Robustness	The auction is unaffected by deviations in the protocol or invalid bids.
Revocation	Bidders can be easily revoked from the CDA if they have broken the rules.

security. We propose an extension to [5] that takes account of the Broker's requirements and responsibilities.

The relationship between a Client and the Broker presents new security problems that have not been previously discussed in auction security literature. One such concern is that a Client needs to know that his/her bid has been passed on to the Auctioneer. For example, a corrupt Broker might deliberately block a Client's bid if they don't want them to initiate a trade. This might happen if the Broker thinks that the Client's judgement is unsound and is acting in a manner inconsistent with the Broker's advice. Alternately, the Broker might have malicious plans for the Client's holdings for a share market related crime such as ramping the market. A further security concern is regarding the Broker's ability to initiate unauthorised trades on a client's behalf. As a Broker knows a Client's personal information, there is nothing preventing them from forging a bid. The bid still appears to be genuine to the Auctioneer.

The Broker and Client share a unique relationship which differs from the Auctioneer/Client relationship. Although the Broker and Client mutually distrust each other, there is some requirement for trust as the parties benefit from cooperation. (i.e., the Client makes a trade and the Broker collects a commission.) These requirements make it difficult to define a clear privacy policy between the Broker and a Client.

As a Broker is providing a service to the Client, operational requirements dictate that a Broker knows certain facts regarding the bids submitted. For example, to determine the Broker's commission and which Client to collect payment from. In addition, the Broker must manage a Client's portfolio and therefore knows everything about the Client's trading behaviour. Therefore, the privacy requirement must be relaxed between a Client and the Broker. Instead, privacy between the Broker and a Client can be enforced with a non-disclosure agreement. Although this is not ideal from an anonymity perspective, it is still an improvement on allowing the Auctioneer and market observers full access to private information.

Figure 8 shows an extension to the registration procedure of [5]. This modification incorporates the Broker's role in the CDA process. The protocol consists of six messages indicated by the arrows. Listed with each arrow is the information sent. A Client initially supplies the Broker with his/her ID. The Broker signs the Client's ID using its private key. The result is $cert_B(ID)$. The Broker sends a Client Number and




Fig. 8. CDA registration protocol incorporating the Broker

$cert_B(ID)$ back to the Client. The rest of the registration procedure is similar to [5].

These modifications allow the Broker to know every bid the Client submits. This is unavoidable as the Client must submit his/her bid via the Broker. However, it does ensure that the Broker can't initiate unauthorised trades, as s/he doesn't know the Client's group certificate. This arrangement allows the Broker to perform the bulk of the work in terms of winner determination, research and portfolio management. A Client need only check the Bulletin Board in the case that they wish to verify that his/her bid has been submitted (i.e., neither the Broker nor the Auctioneer is blocking bids). Furthermore, $cert_B(ID)$ allows the Broker to be associated with a bid during tracing.

VI. CONCLUSIONS

Online share trading is similar to online auctions. As a result, online share trading exhibits similar privacy and security problems. However, there are more parties involved in the share trading process. This makes defining privacy and security requirements more difficult. Once the pretence of corporate provided security is dropped, the bottom line is parties cannot trust each other. A Client might repudiate having made a bid, or forge bids. Likewise, Brokers may steal payments and execute bogus orders. Existing payment enforcement schemes are restrictive and only provide ad-hoc security. Furthermore, the Auctioneer (i.e., Share Market) can match bids in a non-verifiable manner. The parties involved in the share trading process do not have strong privacy policies. An individual's

information can be sold to marketing agencies, or benevolently used against them.

Share market crimes further complicate security matters for CDAs. Insider trading involves using inside knowledge of market sensitive information as the motivation for trading in a particular manner. Misinformation involves deliberately misleading others such that they make an unsound investment, or invest in such a manner that fraudulently benefits the misleader. Existing security mechanisms employed by online Brokers are insufficient and privacy is limited. Regulatory Authorities govern participants and can punish those who engage in criminal or devious behaviour. However, such authorities are expensive to seek recourse through and therefore are largely restricted to policing large-scale fraud. Furthermore, the existence of such an entity also requires the implicit trust of all share market participants.

Cryptographic CDA schemes can prevent many of discussed problems. We show how an existing CDA scheme can be used for secure and private online share trading. Bids are anonymous and the auction process is publicly verifiable. Trust is split between two independent companies, rather than a single trusted Regulatory Authority.

In 1984, David Chaum invented a Digital Cash protocol that allows an individual to spend anonymous cash in a secure manner (see [1]). Security in auction schemes was not significantly addressed until Franklin and Reiter in 1996 (see [2]). Cryptographic share market protocols are a logical extension to digital cash. Combining cash-less protocols with auction schemes to construct a secure and anonymous share market, is perhaps the last step in a truly cash-less society.

REFERENCES

- [1] D. Chaum, "Security without identification: transaction systems to make Big Brother obsolete", *Communications of the ACM*, vol. 29, issue 10, 1985.
- [2] M. Franklin and M. Reiter, "The Design and Implementation of a Secure Auction Service", *IEEE Transactions on Software Engineering*, vol. 22, 1996, pp. 302–312.
- [3] J. Trevathan, "Security, Anonymity and Trust in Electronic Auctions", *ACM Crossroads*, volume 11.3, 2005, pp. 3–9.
- [4] J. Trevathan, H. Ghodosi, and W. Read, "Design Issues for Electronic Auctions", *Proc. 2nd International Conference on E-Business and Telecommunication Networks - ICETE*, pages 340–347, 2005.
- [5] J. Trevathan, H. Ghodosi, and W. Read, "An Anonymous and Secure Continuous Double Auction Scheme", *Proc. 39th International Hawaii Conference on System Sciences - HICSS*, 2006, pp. 125 (1–12).
- [7] J. Trevathan and W. Read, "Secure Online English Auctions", *Proc. International Conference on Security and Cryptography - SECRYPT*, 2006, pp. 387–396.
- [8] K. Viswanathan, C. Boyd and E. Dawson, "A Three Phased Schema for Sealed Bid Auction System Design", *Proc. Australasian Conference on Information Security and Privacy - ACISP*, vol. 1841 of *Lecture Notes in Computer Science*, Springer-Verlag, 2000, pp. 412–426.
- [9] C. Wang and H. Leung, "Anonymity and Security in Continuous Double Auctions for Internet Retail Market", *Proc. 37th Hawaii International Conference on System Sciences - HICSS*, 2004.

7.2 An Anonymous and Secure Continuous Double Auction Scheme

A Continuous Double Auction (CDA) allows many buyers and sellers to continuously submit bids for the purchase and sale of a commodity (e.g., online share trading). Protocols protecting privacy in this type of powerful market mechanism are essential. However, until recently CDA security has been given limited coverage. This paper describes a new scheme for conducting an anonymous and secure CDA. It shows that any existing secure group signature scheme can be used to implement a CDA, which has the following characteristics: unforgeability, anonymity, unlinkability, exculpability, coalition-resistance, verifiability, robustness and traceability. Furthermore, bidders can be added to and removed from the auction without affecting the process of the auction. The proposed scheme is more flexible than the only existing secure CDA scheme, which in contrast provides only a limited subset of these characteristics.

“An Anonymous and Secure Continuous Double Auction Scheme” [5] is published in the proceedings of the Thirty-Ninth Hawaii International Conference on System Sciences (HICSS’06). HICSS’06 was held at the Hyatt Regency Kauai Resort and Spa on the Island on Kauai. Sponsors included the University of Hawaii and ACM. The proceedings consist of over 450 papers in nine major tracks and two symposia. This paper was presented as part of the mini track: Information Systems Security Management. 11 papers were submitted to this track, which were subjected to a rigorous reviewing process. In all, 6 papers were accepted with an acceptance rate close to 50%. This paper is available online via ACM Portal ¹, and is referenced by DBLP ².

This paper is co-authored by Dr Hossein Ghodosi and received some initial comments from Associate Professor Bruce Litow. Dr Ghodosi’s research interests include Cryptography and Information Security. Publication in HICSS was made possible via the Graduate Research Scholarship (GRS) and the School of Mathematical and Physical Sciences.

¹<http://www.portal.acm.org>

²<http://www.informatik.uni-trier.de/~ley/db/>

An Anonymous and Secure Continuous Double Auction Scheme

Jarrod Trevathan
School of Mathematical and
Physical Sciences
James Cook University
Email: jarrod.trevathan@jcu.edu.au

Hossien Ghodosi
School of Information Technology
James Cook University
Email: hossein@cs.jcu.edu.au

Wayne Read
School of Mathematical and
Physical Sciences
James Cook University
Email: wayne.read@jcu.edu.au

Abstract

A Continuous Double Auction (CDA) allows many buyers and sellers to continuously submit bids for the purchase and sale of a commodity (e.g., online share trading). Protocols protecting privacy in this type of powerful market mechanism are essential. However, until recently the security of CDAs has been given limited coverage. This paper describes a new scheme for conducting an anonymous and secure CDA. We show that any existing secure group signature scheme can be used to implement a CDA which has the following characteristics: unforgeability, anonymity, unlinkability, exculpability, coalition-resistance, verifiability, robustness and traceability. Furthermore, bidders can be added to and removed from the auction without affecting the process of the auction. Our scheme is more flexible than the only existing secure CDA scheme, which in contrast provides only a limited subset of these characteristics.

I. Introduction

An auction is an exchange mechanism where many potential buyers submit bids for a commodity, which is usually awarded to the highest bidder. The Auctioneer accepts bids on behalf of the seller of the commodity and determines the winner according to the auction rules.

An auction that has only one seller and many buyers is referred to as a *single* auction. A *Continuous Double Auction (CDA)* allows for many buyers and sellers to continuously submit bids for the purchase and sale of a commodity. The most well known use for CDAs is in share markets.

In recent years, several companies have emerged that offer auctioning services via an online network such as eBay¹ and OnSale². These types of auctions have geographical advantages over traditional auctions, as sellers and buyers need not be physically present at a central location during the auction proceedings. This allows for larger and more elaborate auctions, which reach many more bidders than traditional auctions. However, it also allows the participants to cheat.

It is well documented that major security problems in online auctions exist (e.g., see [7], [17], [18], [20]). For example, buyers and sellers might be in collusion, repudiate bids, or not pay for/deliver the item bid upon. Alternately, the Auctioneer might be corrupt and award the auction to someone other than the legitimate winner. A further significant issue is how to protect a bidder's identity and bidding information.

An emerging aspect of electronic auctions is the ability to trade shares on the Internet. This is done using a brokerage

site such as Commsec³. A broker submits buy and sell bids to the Auctioneer running a CDA. Present brokers only offer trivial security measures such as issuing a client with a user name and password, and encrypting trading sessions. During this process much information about a client's identity and bidding pattern is released to the public. In terms of a share market, certain information is required by law to be available to authorised parties. Beyond this, it is the right of the bidder to have his/her anonymity protected.

Privacy is a desirable requirement in online trading for numerous reasons. Firstly, it prevents the Auctioneer from selectively blocking bids, based on the identity of a bidder. Secondly, many banking institutions tend to sell client's information to advertising agencies. Finally, share trading can be much more controversial than normal auctioning. For example, there may be legal/industrial disputes, bankruptcies, criminal behaviour (insider trading), etc. In such cases, there must be measures to protect an individual's anonymity (or revoke it when required).

Coverage exists regarding different CDA models (see e.g., [8], [9]) and bidding strategies (see e.g., [10], [16]). However, security and anonymity in CDAs has been given limited attention. Wang and Leung's scheme [21], to our knowledge, is the only auction scheme proposed that specifically addresses CDA security. In their CDA scheme, trust is divided amongst two separate servers in order to protect a bidder's anonymity. However, we have discovered that this scheme allows the identity of a bidder to be revealed immediately after his/her first bid. Furthermore, it allows profiles to be created about a bidder's trading behaviour, as bids are linkable. In this paper, we propose a secure CDA scheme to fix these problems.

This paper is organised as follows: In the remainder of this section we describe the fundamentals of electronic auctions, general electronic auction security and our contribution to CDA security. In Section II, we will review Wang and Leung's scheme, and describe issues specific to CDAs. In Section III, we will discuss the components of our new CDA scheme including the communication model, parties involved in the system, and cryptographic building blocks employed in our scheme. Section IV outlines the protocol for our CDA. In Section V, we will discuss security issues for the proposed CDA. Section VI gives an efficiency comparison of Wang and Leung's scheme to several different implementations of our

¹<http://www.ebay.com>

²<http://www.onsale.com>

³<http://www.commsec.com.au>

scheme.

A. *Electronic Auction Fundamentals and Security Issues*

The main activities fundamental to any electronic auction are:

Initialisation – The Auctioneer sets up the auction and advertises it i.e., type of good being auctioned, starting time, etc.

Registration – In order to participate in the auction, bidders must first register with the Auctioneer. This ensures that only valid bids are made, and that bidders can be identified for payment purposes.

Bidding – A registered bidder computes his/her bid and submits it to the Auctioneer. The Auctioneer checks the bid received to ensure that it conforms with the auction rules.

Winner Determination – The Auctioneer determines the winner according to the auction rules. It is desirable that the determination of the auction winner can be publicly verified.

The security requirements for an electronic auction are numerous. In general, core requirements include that the Auctioneer must determine the correct winner according to the auction rules; the winner must receive the item from the Seller; the Seller must receive payment in full from the winning bidder; and no bidder should have more information than any other to determine his/her bid. The final requirement is especially true in CDAs, where insider trading is considered an offence. Discussions regarding auction security can be found in [3], [7], [13], [20]. The following outlines the general security goals for electronic auctions:

Unforgeability – Bids must be unforgeable, otherwise a bidder can be impersonated.

Non-Repudiation – Once a bidder has submitted a bid, they must not be able to repudiate having made the relevant bid. Otherwise, if a bidder wins and does not want to pay they might deny that they submitted the winning bid.

Public Verifiability – There must be publicly available information by which all parties can be verified as having correctly followed the auction protocol. This should include evidence of registration, bidding and proof of the winner of the auction.

Robustness – The auction process must not be affected by invalid bids nor by participants not following the correct auction protocol.

Anonymity – The bidder-bid relationship must be concealed so that no bidder can be associated or identified with the bid they submit.

Auctions can be classified according to whether the bids are *sealed* or *open*. In a sealed bid auction, bidders submit their bids during a bidding round. The bid values remain sealed (i.e., no one else knows what the bid values are) until a winner determination round. Open bid auctions on the other hand allow everyone to know the value of everyone else's bids. Sealed bid auctions introduce their own unique security concerns which have been well-documented in existing literature (see [7], [13], [20]). However, CDAs are open bid, therefore no one gains any advantage by knowing the values of other people's bids. In

fact it is essential for CDAs to provide information regarding market depth (i.e., list of buy/sell bids and their values) to the participants, in order to facilitate the bidding process.

B. *New CDA Results and our Contribution*

In this paper we present a new proposal for conducting anonymous and secure CDAs. We show that the CDA scheme by Wang and Leung is inadequate and describe several concerns regarding its security and anonymity claims. Furthermore, we explain how any existing secure group signature scheme can be used to provide anonymity and security for a CDA.

Chaum and van Heyst [6] introduced the concept of group signatures. A group signature scheme allows members of a group to sign messages on behalf of the group, such that the resulting signature does not reveal their identity. Signatures can be verified with respect to a single group public key, while hiding the identity of the signer. Only a designated group manager is able to open signatures, and thus reveal the signer's identity.

Since its invention, group signatures have been the subject of investigation by many authors. Camenisch and Stadler [5] present the first efficient group signature scheme, where the size of the group's public key, and the signatures, are independent from the number of group members. Furthermore, the group's public key remains unchanged if a new member is added to the group. Group signatures have desirable properties which makes them ideal for use in electronic auctions.

By incorporating a group signature scheme, our CDA provides the following characteristics:

Unforgeability – Only registered bidders can issue a bid.

Anonymity – The identity of a bidder remains anonymous (i.e., a bidder cannot be associated with the bid s/he has submitted). However, the value of the bid is still available to others (i.e., open bid).

Unlinkability – No party can form a profile about a bidder's bidding strategy based on past bids.

Traceability – In case of dispute and/or to determine the winner of the auction, the authorities (e.g., the Registrar and Auctioneer, cooperatively) are always able to identify the bidder.

Exculpability – No party –even the Registrar and Auctioneer– can issue a bid on behalf of any other bidder.

Coalition-resistance – A colluding subset of registered bidders cannot issue a bid that cannot be linked back to one of the colluding bidders.

Verifiability – All auction participants can be verified as having followed the protocol correctly.

Robustness – Invalid bids or failures in following the protocol does not affect the overall auction outcome.

Revocation – Bidders can be easily revoked from the CDA if they have broken the rules.

Note that our CDA scheme has more security characteristics than those outlined in Section I-A (i.e., unlinkability, traceability, exculpability, coalition-resistance and revocation). CDAs are quite unique compared to more conventional types

of auctions and the reasons for these additional security requirements will become apparent once we examine an existing secure CDA proposal.

II. Existing CDA Scheme

Wang and Leung's scheme [21] is the first attempt to conduct secure CDAs while preserving the anonymity of bidders. In Wang and Leung's CDA, there are two authorities, the Registration Manager (RM) and the Market Manager (MM). Before entering the CDA market, each bidder must register with both RM and MM. RM holds the corresponding relation of the identity and the certificate of bidders. MM verifies the certificate and issues a pseudonym for the bidder.

Registration with RM in [21] -

Suppose RM uses the RSA system with public parameters n, e as the RSA modulo and the encryption key, respectively. Let d be the corresponding private key in the RSA system and $H(\cdot)$ be an appropriate hash function. Denote by $\Omega_x(m)$ the set $\{m, \sigma_x(H(m))\}$, that is a message, m , and the signature of party x on $H(m)$. Also, let a bidder, b_i , possess a certified encryption/decryption key pair. The registration protocol between b_i and RM is shown in Table I (note that \parallel denotes concatenation).

After registration with RM, the bidder must register with MM. Table II illustrates the protocol between MM and bidders.

Upon obtaining the certificate from MM, the bidder b_i submits their bid in the form $\{\text{offer}, \sigma_{b_i}(\text{offer}), \text{Cert}_{b_i}\}$, where $\text{offer} = \{p_{b_i}, \text{Buy/Sell}, \text{Commodity}, \text{Value}, \text{Timestamp}\}$, and p_{b_i} is a temporary public-key (or the pseudonym) associated with bidder b_i .

A. Analysis of Wang-Leung's Scheme

- 1) During registration with MM, the bidder forms the message M_4 , which contains the message M_2 as the first component. However, M_2 is formed as $\Omega_{RM}(ID_{b_i} \parallel sn)$ (i.e., it contains the identity of bidder b_i). Therefore, MM immediately learns the true identity of b_i .
- 2) Bids are linkable in this scheme, as the same pseudonym is used for each bid. This allows profiles to be created about bidders.
- 3) It is unclear how this scheme identifies winners. After the first trade (i.e., the first time a bidder wins) the identity of the bidder (associated with its pseudonym) is disclosed, and thus, later bids issued by this bidder are not anonymous.
- 4) When a bidder's identity has been traced, this scheme reveals the values of all past bids they have made. This is undesirable if there is no reason to doubt the validity of the guilty bidder's previous bids.

B. Further Problems - Procrastinating Attack

Even if the protocol by Wang and Leung is fixed, anonymity cannot be achieved in this model. In this section we will propose a procrastinating attack that can be used to immediately reveal a bidder's true identity after registration.

However, before we outline the attack, we must first examine the basic components of an auction. In a traditional auction model, there are two parties (excluding the Seller):

- 1) The Auctioneer, who organises the auction, accepts the bids on behalf of the seller of the commodity, and determines the winner according to the auction's rules.
- 2) The Bidder, who computes his/her bid and submits it to the Auctioneer, attempting to win the commodity according to the auction's rule.

This simple model works based on the assumption that the Auctioneer is trustworthy. However, the major problem in electronic auction protocols is to protect bidders from a corrupt Auctioneer. Franklin and Reiter [7] suggest distributing the role of Auctioneer amongst ℓ servers. The auction can be considered secure/fair unless a threshold t , $1 \leq t \leq \ell$ of the servers collude. However, these types of schemes have a high communication overhead and cannot be trusted when the same company owns all the servers.

An alternative approach is to split trust among two servers owned by separate entities (see [11]–[13]). The auction is considered secure as long as the two parties do not collude. Let S_1 and S_2 denote the two servers/parties. This model essentially works as follows;

- 1) Bidders present their identification to S_1 and obtain a token that does not carry their ID.
- 2) Bidders submit the token (without revealing their ID) to S_2 , which issues a pseudonym associated with the token.

In this way, neither S_1 nor S_2 know the relationship between any real ID and the pseudonyms. The bidder then submits his/her bid using the pseudonym. In case of dispute, however, S_1 and S_2 (cooperatively) can determine the real ID associated with each pseudonym. This technique works properly for single auctions, in which the registration phase is separate from the bidding phase. However, in CDAs, the registration and bidding phases overlap and this protocol is not secure. The following scenario explains our attack:

The Attack – A bidder provides identification to S_1 and obtains a token. Using this token, it gets a pseudonym from S_2 that can be used for bidding. If there is no separation between the registration and bidding phases, S_1 can act in a procrastinating manner by halting all new registrations until the recently registered bidder submits his/her bid. This enables S_1 to learn the mapping between the bidder's identity and his/her pseudonym. As a result, Wang and Leung's CDA [21] scheme is susceptible to this attack.

Note that Wang and Leung's scheme is given in the context of Internet retail markets whereas our scheme is designed specifically for share market applications. There are differing security and anonymity requirements between these two types of CDA. For example, a share market offers identical intangible goods and the valuation (i.e., the price) is calculated in a different manner (i.e., there is no "wear and tear" to depreciate its value, instead the value is calculated using a company's

TABLE I
THE REGISTRATION PROTOCOL BETWEEN b_i AND RM

Bidder		RM
Generates a random sn , and $M1 = \Omega_{b_i}(ID_{b_i}, request, sn)$	$\xrightarrow{M1}$	Verifies b_i 's signature Generates $M2$, such that: $M2 = \Omega_{RM}(ID_{b_i} sn)$
	$\xleftarrow{M2}$	
Verifies RM's signature Generates random numbers r Computes pseudonym $ps_{b_i} = H(H(ID_{b_i} sn r))$ Generates a pair of temporary private/public keys sb_i, pb_i	$\xleftrightarrow{\text{Cut-and-Choose Protocol}}$	
Obtains $\sigma_{RM}(ps_{b_i}, pb_i)$		

TABLE II
THE REGISTRATION PROTOCOL BETWEEN b_i AND MM

Bidder		MM
Generates M4 as: $\{M2, r, (ps_{b_i}, pb_i), \sigma_{RM}(ps_{b_i}, pb_i)\}$	$\xrightarrow{M4}$	Make sure that b_i has not been registered before, If M4 is verified, then generate $Cert_{b_i}$ as: $\{\Omega_{RM}(ps_{b_i}, pb_i), \sigma_{MM}(ps_{b_i}, pb_i)\}$
	$\xleftarrow{Cert_{b_i}}$	
Obtains the auction certificate.		

expected profit). Furthermore, the location of the seller in a share market has no effect on the cost of shipping the item as occurs in auctions with more conventional items. Finally, it is unclear to the authors exactly what an Internet retail market is as eBay (and similar online auction sites) do not actually list this as an auction type that buyers/sellers can use. Share markets on the other hand are the best known example of a CDA.

III. Components of our CDA Scheme

The scheme we present in this paper uses the additional security characteristics of unlinkability and revocation to fix the problems of the existing scheme. As our CDA is based on a group signature, coalition resistance and exculpability are included to ensure that neither group members nor those involved in issuing a member's group certificate can forge an innocent bidder's signature. Note that non-repudiation is referred to as "traceability" in our CDA scheme.

A. Parties in The System

Figure 1 illustrates the major parties, stages and information flow of our CDA scheme. The auction has three parties:

- 1) A Bidder, who is interested in buying and/or selling commodities in a CDA.
- 2) An Auctioneer, who organises the auction, accepts the bids and determines the winner according to the auction

rules. The Auctioneer also holds the corresponding relation of the identity and a token associated with each bidder.

- 3) A Registrar, who takes part in a protocol in order to complete the registration of a bidder who has obtained a token from the Auctioneer. At the end of the protocol, the bidder obtains a secret key that enables him/her to generate signed bids in a proper format.

The arrows in Figure 1 represent the messages sent between parties. Dashed circles surrounding messages indicate that the enclosed messages are sent during a particular stage of the auction protocol. Each stage of the protocol is explained in Section IV.

B. Communication Channel

Each bidder and the Auctioneer are connected to a common broadcast medium (or public Bulletin Board) with the property that messages sent to the channel instantly reach every party connected to it. We assume that the broadcast channel is public, that is, everybody can listen to all information communicated via the channel, but cannot modify it. We further assume that there are also private channels between the Registrar and any potential bidders (who wish to join the CDA). Nobody is able to listen to, or modify, the messages




Fig. 1. System Dynamics of the CDA Scheme

sent via these private channels ⁴.

C. Group Signatures

For the CDA construction presented in the paper, we employ a group signature scheme invented by Ateniese *et al.* [1], which is provably secure. To join a CDA, a bidder must first register with a Registrar (who jointly plays the role of a group manager with the Auctioneer in a group signature scheme). Once registered, a bidder can participate in the CDA by signing bids using the group signature. Bids are submitted to an independent Auctioneer who runs the CDA. The auction results are posted on a publicly verifiable bulletin board by the Auctioneer. The Ateniese *et al.* [1] group signature scheme informally works as follows:

Let $n = pq$ be an RSA modulus, where p and q are two safe primes (i.e., $p = 2p' + 1$, $q = 2q' + 1$, and p' , q' are also prime numbers). Denote by $QR(n)$, the set of quadratic residues – a cyclic group generated by an element of order $p'q'$. The group public key is $\mathcal{Y} = (n, a, a_0, y = g^x, g, h)$, where a, a_0, g, h are randomly selected elements from $QR(n)$. The secret key of the group manager is x .

To join the group, a user (bidder b_i) must engage in a protocol with the group manager (Registrar) to receive a group certificate $[B_i, e_i]$, where $B_i = (a^{x_i} a_0)^{1/e_i} \bmod n$ with e_i and x_i chosen from two integral ranges as defined in [1] (x_i is only known to the user/bidder).

In order to sign a message/bid, m , the user/bidder has to prove possession of his member certificate $[B_i, e_i]$ without revealing the certificate itself. More precisely, the user/bidder computes:

$$T_1 = B_i y^w \bmod n, \quad T_2 = g^w \bmod n, \\ T_3 = g^{e_i} h^w \bmod n, \quad SK(m)$$

⁴This assumption can be realised by using cryptographic tools that provide authenticated communication.

where the value $SK(m)$, computed over a message m , indicates a signature of knowledge of the secret key x_i and the e_i th root of the first part of the representation of T_3 (in the implementation of our scheme, the exact signature generation and verification procedures will be presented).

In case of a dispute, the group manager can open a signature that reveals the identity of the signer. This is due to the fact that the pair (T_1, T_2) is an ElGamal encryption of the user's certificate (using the public key of the group manager). That is, the group manager can compute B_i , using $B_i = T_1 / (T_2)^x$.

IV. CDA Protocol

In this section we present the protocol for our CDA scheme. Note that unlike single auctions, CDAs do not consist of various separate phases, as the auction runs continuously and thus all phases overlap. Some bidders are in the process of buying/selling while others may concurrently be in the process of registration. However, from the point of view of an individual bidder and/or the nature of activities in the auction, our scheme consists of the stages below.

A. Setup

Most activities of this stage need to be performed only once (in order to establish a CDA). The Auctioneer organises the CDA (i.e., advertising and calls for auction). Let $\lambda_1, \lambda_2, \xi_1$, and ξ_2 be some lengths, Λ, Γ be some integral ranges, and $\mathcal{H}(\cdot)$ be a collision-resistant hash function. The Registrar sets up the group public key and his secret key by performing the following steps:

- 1) Chooses two safe primes p and q (i.e., $p = 2p' + 1$ and $q = 2q' + 1$, where p', q' are prime numbers) and sets the RSA modulus $n = pq$.
- 2) Chooses random elements $a, a_0, g, h \in QR(n)$.
- 3) Chooses a secret element $x \in_R Z_{p'q'}^*$, and sets

$$y = g^x \pmod{n}$$

4) Publishes the group's public key as

$$\mathcal{Y} = (a, a_0, g, h, n, y)$$

B. Registration

A user submits a request to the Auctioneer to participate in the CDA. The Auctioneer verifies the identity of the requestor, and issues a token that is verifiable by the Registrar. The user then takes part in a protocol with the Registrar to obtain his secret key and a certificate of membership in the auction. Note that the token does not carry the real identity of the bidder. All communication between the Registrar and the owner of a token is authenticated and kept for tracing purposes (i.e., revealing the identity of user associated with the token).

Suppose the Auctioneer uses the RSA system with public parameters n, e as the RSA modulo and the encryption key, respectively. Let d be the corresponding private key in the RSA system and $H(\cdot)$ be an appropriate hash function. Denote by $\Omega_x(m)$ the set $\{m, \sigma_x(H(m))\}$, that is a message, m , and the signature of party x on $H(m)$. Also, let a bidder, b_i , possess a certified encryption/decryption key pair. The protocol between a new user and the Auctioneer is as follows:

- 1) The user establishes his/her identity with the Auctioning by sending a *request* to participate in the CDA, identification, ID_{b_i} , and a random number, sn . This is signed using b_i 's private key, σ_{b_i} .
- 2) The Auctioneer verifies the user's signature using b_i 's public key. The Auctioneer stores all of the information and signs sn using the private key σ_A . The Auctioneer stores the token, $\Omega_A(m)$, securely and forwards a copy to b_i .

This process is shown in Table III. The Auctioneer's signature on the user's identity, $M2$ (see Table III), serves as a token that the user can present to the Registrar. The Registrar is able to check that it is a valid token issued by the Auctioneer by using the Auctioneer's public key.

The protocol between a new bidder, b_i , and the Registrar is as follows (checks in which the values are chosen from proper intervals, the user knows discrete logarithms of values, etc. are omitted):

- 1) b_i selects random exponents x'_i, r and sends

$$C_1 = g^{x'_i} h^r \pmod{n}$$

to the Registrar.

- 2) The Registrar checks that $C_1 \in QR(n)$. If this is the case, the Registrar selects random values

$$\alpha_i, \beta_i$$

and sends them to b_i .

- 3) b_i computes

$$x_i = 2^{\lambda_1} + (\alpha_i x'_i + \beta_i \pmod{2^{\lambda_2}})$$

and sends to the Registrar the value

$$C_2 = a^{x_i} \pmod{n}$$

- 4) The Registrar checks that $C_2 \in QR(n)$. If this is the case, the Registrar selects a random $e_i \in \Gamma$ and computes

$$B_i = (C_2 a_0)^{1/e_i} \pmod{n}$$

and sends to b_i the membership certificate

$$[B_i, e_i]$$

(note that $B_i = (a^{x_i} a_0)^{1/e_i} \pmod{n}$).

- 5) b_i verifies that

$$a^{x_i} a_0 = B_i^{e_i} \pmod{n}$$

- 6) The Registrar publishes value f , where

$$f := e_1 * e_2 * \dots * e_\ell$$

(e_i s, $1 \leq i \leq \ell$, are the exponents of current group members).

The Registrar then creates a new entry in the membership table and stores b_i 's membership certificate $[B_i, e_i]$ and a transcript of the registration process in this location.

C. Bidding

Using a membership certificate $[B_i, e_i]$, a bidder can generate anonymous and unlinkable group signatures on a bid. A bidder b_i , submits a signed bid, to the Auctioneer of the form:

$$\{Buy/Sell, Commodity, Value, Timestamp, etc.\}$$

In securities markets, it is typical for a bid to also contain information regarding the quantity being bid for (e.g., a bidder desires 20 units of commodity at a price of \$10 per unit). In addition, an expiration time for the bid might also be included indicating that a bid is valid until a particular date, after which it should be discarded from the auction proceedings. Furthermore, conventional auctions usually allow a bidder to submit only one bid, whereas CDAs allow an indefinite number of bids to be submitted (e.g., a bidder may buy and sell the same commodity, offering different prices based on the quantity of the goods, etc.) In order to generate a signature on a message/bid, m , a bidder b_i performs the following:

- 1) Chooses a random value w and computes:

$$T_1 = B_i y^w \pmod{n}, \quad T_2 = g^w \pmod{n},$$

$$T_3 = g^{e_i} h^w \pmod{n}$$

- 2) Chooses r_1, r_2, r_3, r_4 (in random) from predetermined intervals and computes:

$$(a) \quad d_1 = T_1^{r_1} / (a^{r_2} y^{r_3}), \quad d_2 = T_2^{r_1} / (g^{r_3}), \quad d_3 = g^{r_4},$$

$$\text{and } d_4 = g^{r_1} h^{r_4} \text{ (all in mod } n),$$

$$(b) \quad c = \mathcal{H}(g \parallel h \parallel y \parallel a_0 \parallel a \parallel T_1 \parallel T_2 \parallel T_3 \parallel d_1 \parallel d_2 \parallel d_3 \parallel d_4 \parallel m),$$

$$(c) \quad s_1 = r_1 - c(e_i - 2^{\xi_1}), \quad s_2 = r_2 - c(x_i - 2^{\lambda_1}),$$

$$s_3 = r_3 - c e_i w, \quad \text{and } s_4 = r_4 - c w \text{ (all in } \mathbb{Z}).$$

- 3) Chooses a random number, win , which is used to verify ownership of the bid during the winner determination stage. This is kept secret by both b_i and the Auctioneer.

TABLE III
THE REGISTRATION PROTOCOL BETWEEN A NEW USER AND THE AUCTIONEER

New User	Auctioneer
Generates a random sn , and $M1 = \Omega_{b_i}(ID_{b_i}, request, sn)$	Verifies the user's signature Stores ID_{b_i} and sn Generates $M2$, such that: $M2 = \Omega_A(sn)$
Verifies the Auctioneer's signature	

TABLE IV
THE REGISTRATION PROTOCOL BETWEEN b_i AND THE REGISTRAR

Bidder	Registrar
Selects random exponents x'_i, r Calculates $C_1 = g^{x'_i h^r} \pmod n$ $M3 = \{M2, C_1\}$	Make sure that b_i has not been registered before. Checks that $C_1 \in QR(n)$, Selects random values α_i, β_i $M4 = \{\alpha_i, \beta_i\}$
Computes $x_i = 2^{\lambda_1} + (\alpha_i x'_i + \beta_i \pmod{2^{\lambda_2}})$ $C_2 = a^{x_i} \pmod n$ $M5 = \{C_2\}$	
Verifies that $a^{x_i} a_0 = B_i^{e_i} \pmod n$	Checks that $C_2 \in QR(n)$ Selects a random $e_i \in \Gamma$, Computes $B_i = (C_2 a_0)^{1/e_i} \pmod n$ $M6 = \{B_i, e_i\}$ Publishes $f := e_1 * e_2 * \dots * e_l$

TABLE V
THE BIDDING PROTOCOL BETWEEN b_i AND THE AUCTIONEER

Bidder	Auctioneer
Chooses a random value w and computes: $T_1 = B_i y^w \pmod n$, $T_2 = g^w \pmod n$, $T_3 = g^{e_i} h^w \pmod n$ Chooses r_1, r_2, r_3, r_4 (in random) from predetermined intervals Computes: $d_1 = T_1^{r_1} / (a^{r_2} y^{r_3})$, $d_2 = T_2^{r_1} / (g^{r_3})$, $d_3 = g^{r_4}$, and $d_4 = g^{r_1} h^{r_4}$ (all in mod n), $c = \mathcal{H}(g \parallel h \parallel y \parallel a_0 \parallel a \parallel T_1 \parallel T_2 \parallel T_3 \parallel d_1 \parallel d_2 \parallel d_3 \parallel d_4 \parallel m)$, $s_1 = r_1 - c(e_i - 2^{\xi_1})$, $s_2 = r_2 - c(x_i - 2^{\lambda_1})$, $s_3 = r_3 - c e_i w$, and $s_4 = r_4 - c w$ (all in \mathbb{Z}) Chooses a random number, win . $M7 = (c, s_1, s_2, s_3, s_4, T_1, T_2, T_3, win)$	Verifies the user's signature Compute (all in mod n): $c' = \mathcal{H}(g \parallel h \parallel y \parallel a_0 \parallel a \parallel T_1 \parallel T_2 \parallel T_3 \parallel$ $(a_0^c T_1^{(s_1 - c 2^{\xi_1})}) / (a^{s_2 - c 2^{\lambda_1}} y^{s_3}) \parallel$ $(T_2^{s_1 - c 2^{\xi_1}}) / (g^{s_3}) \parallel T_2^c g^{s_4} \parallel T_3^c g^{s_1 - c 2^{\xi_1}} h^{s_4} \parallel m)$ Accepts the signature iff $c = c'$ and the parameters s_1, s_2, s_3, s_4 lie in proper intervals Posts the bid on the Bulletin Board
Observes the Bulletin Board	

4) The bid submitted to the Auctioneer is
 $(c, s_1, s_2, s_3, s_4, T_1, T_2, T_3, win)$

The Auctioneer then checks the validity of the bidder's signature on bid, using the group's public key \mathcal{Y} . A bid of

the correct form is included in the CDA (i.e., posted on the bulletin board). An invalid bid is discarded. The Auctioneer verifies the signature as follows:

- 1) Compute (all in mod n):

$$c' = \mathcal{H}(g \parallel h \parallel y \parallel a_0 \parallel a \parallel T_1 \parallel T_2 \parallel T_3 \parallel (a_0^c T_1^{(s_1 - c2^{\xi_1})}) / (a^{s_2 - c2^{\lambda_1}} y^{s_3}) \parallel (T_2^{s_1 - c2^{\xi_1}}) / (g^{s_3}) \parallel T_2^c g^{s_4} \parallel T_3^c g^{s_1 - c2^{\xi_1}} h^{s_4} \parallel m).$$
- 2) Accept the signature iff $c = c'$ and the parameters s_1, s_2, s_3, s_4 lie in proper intervals.

D. Winner Determination

The Auctioneer determines the auction outcome according to general CDA rules. Conventional auctions only allow for one (or few) winner(s), however, CDAs can have multiple winners. This results in more complicated auction rules.

Bids are referred to as *cleared* when the value of a buy bid equals the value of a sell bid (e.g., if there is a buy bid of \$2 and a sell bid of \$2, then the Auctioneer can match and clear these two bids). The Australian Stock Exchange ⁵ clears bids in a strictly time based manner. No priority is given to bids based on amount. Matching buy and sell bids are cleared. In the situation where a buy bid exceeds a sell bid, the surplus (i.e., the difference between the buy and sell price) is awarded to the seller.

Sandholm *et al.* [15] examines various techniques for performing online market clearing. They show that the Auctioneer can increase the market volume (number of cleared bids) by retaining the proceeds of surplus trades. Surplus revenue can then be used to subsidise loss-making trades.

Wurman *et al* present an academic auction server referred to as “AuctionBot” in [22]. The purpose of this server is to allow researchers to explore the design space of auctions. In [23] Wurman and Wellman describe several CDA clearing algorithms that have been tested on AuctionBot.

In our scheme, as soon as bids have been verified, they are posted to the bulletin board in plaintext. This allows the Auctioneer to employ any auction rules or clearing algorithm to the bids. This method has been used by Viswanathan *et al* [20], and Wang and Leung [21] to provide verifiability of the matching process. A full description of clearing strategies is orthogonal to our discussion on CDA security.

Once a bid has been cleared, the winner determination protocol proceeds as follows:

- 1) The Auctioneer signs a cleared bid using its private key σ_A . The signed bid and associated information regarding the trade is posted on the bulletin board.
- 2) b_i views the bulletin board to see if he/she has won. b_i presents the winning bid and the random number *win* chosen during the bidding stage to the Auctioneer.
- 3) The Auctioneer checks that *win* matches the winning bid and awards the goods to b_i .

The purpose of *win* is to ensure only the winner is awarded the goods as no one else knows the correspondence between the winning bid and *win*.

⁵<http://www.asx.com.au>

E. Traceability

In the event of a dispute, the Registrar can open the signature on a bid to reveal the certificate assigned to the bidder. This process is as follows:

- 1) Check the signature’s validity via the verification procedure.
- 2) Recover B_i as

$$B_i = T_1 / (T_2)^x \pmod{n}$$

The Registrar then checks the registration transcripts, and determines the token associated with this certificate. The Auctioneer, who knows the relation between tokens and real identities, can determine the identity of the bidder. Note that in our scheme, revealing the identity of a bidder does not reveal any information about his/her past bids.

F. Revocation

In the case of a malicious bidder (i.e., a bidder that has been caught breaking the auction rules) we would like an efficient and secure means of revoking the bidder from the group. In our scheme we employ a technique presented in Ateniese, Song and Tsudik [2], which considers the revocation problem in Ateniese *et al* [1]. The revocation algorithm informally works as follows:

Assume the group consists of ℓ members with associated certificates $[B_1, e_1], \dots, [B_\ell, e_\ell]$ and the value $f = e_1 * e_2 * \dots * e_\ell$ is known to all members. Suppose a member with exponent e_k , ($1 \leq k \leq \ell$) needs to be revoked from the system. The Registrar chooses a random value $u \in QR(n)$ and computes

$$t = u^{1/(f/e_k)} \pmod{n}$$

Then the Registrar publishes t, e_k and the new public key

$$\mathcal{Y} = (a, a'_0, g, h, n, y)$$

where $a'_0 = a_0 * u$. A bidder with exponent e_i , $i \neq k$, updates his group signing certificate as

$$B'_i = B_i * t^{s_i}$$

where $s_i = f / (e_i * e_k)$. So the new certificate for this user is $[B'_i, e_i]$, where

$$B'_i = (a^{x_i} a'_0)^{1/e_i} \pmod{n}$$

Now bidder k , cannot compute the corresponding B'_k and thus cannot sign/issue any new bid.

V. Security

In this section we will discuss how our scheme has addressed the security criteria we proposed for CDAs in the introductory section.

Unforgeability – Only bidders that are members of the group are able to generate a group signature, and thus issue a bid. This is due to the unforgeability of the underlying group signature. Furthermore, each bid contains a timestamp that

TABLE VI
COMPARISON OF CDA SCHEMES

	WL04	ACJT00	CL02	TX03
Registration	$L \times 3 + 2$ exp.	30 exp.	30 exp.	2 exp.
Signing	1 exp.	25 exp.	25 exp.	17 exp.
Verification	3 exp.	21 exp.	21 exp.	16 exp.
Revocation	Unclear	$O(\ell)$	$O(1)$	$O(1)$

serves as a freshness indicator. If an opponent attempts to replay a bid with an outdated timestamp, it will be rejected.

Anonymity – Given a valid signature $(c, s_1, s_2, s_3, s_4, T_1, T_2, T_3)$ identifying the actual signer is computationally difficult. Determining which bidder with certificate $[B_i, e_i]$ has signed a bid, requires deciding whether the three discrete logarithms $\log_y T_1/B_i$, $\log_g T_2$, and $\log_g T_3/(g^{e_i})$ are equal. This is assumed to be infeasible under the Decisional Diffie-Hellman Problem, and thus anonymity is guaranteed. Note that in our CDA, the Registrar can figure out the certificate associated with each signature, but cannot determine the identity of the bidder associated with this certificate.

Unlinkability – Given two signatures $(c, s_1, s_2, s_3, s_4, T_1, T_2, T_3)$ and $(c', s'_1, s'_2, s'_3, s'_4, T'_1, T'_2, T'_3)$, deciding if they are computed by the same bidder is computationally difficult (with the same argument as anonymity).

Exculpability – Neither a bidder nor the Auctioneer and/or Registrar (individually or cooperatively) can sign a message/bid on behalf of another bidder. This is because the secret key of x_i , associated to user i is computationally hidden from the Registrar. The Registrar, at most, can learn $a^{x_i} \bmod n$, which cannot help him/her to learn the exponent x_i (since the discrete logarithm over the safe composite modulo, n , is difficult).

Coalition resistance – This is due to the following theorem.

Theorem 1: [1] Under the strong RSA assumption, a group certificate $[B_i = (a^{x_i} a_0)^{1/e_i} \bmod n, e_i]$ with $x_i \in \Lambda$ and $e_i \in \Gamma$ can be generated only by the group manager provided that the number K of certificates the group manager issues is polynomially bounded.

Verifiability – All bids/signatures are posted to the public bulletin, therefore all parties can verify the auction outcome. A bid can be verified using the group's public key \mathcal{Y} . The winning bidder can present *win* as proof that he/she won.

Robustness – Invalid bids will not be posted on the bulletin board. Moreover, malicious bidders will be revoked from the system, and thus cannot affect the auction outcome.

Traceability – The Registrar is always able to open a valid signature and, with the help of Auctioneer, identify the signer of the bid.

Revocation – Bidders can be easily revoked from the CDA if they have broken the auction rules. This is done by updating the group's public key and each remaining bidder's certificate.

A. Procrastinating Attack

Our scheme is not susceptible to the procrastinating attack described in Section II-B. To see how let the Auctioneer be S_1 and the Registrar S_2 . S_1 attempts to learn the identity of b_i by delaying all new registrations until b_i submits his/her first bid. In our scheme S_1 cannot learn any information about b_i as bids are not linkable (due to the properties of the group signature). When b_i submits a bid, S_1 cannot be certain if the signature belongs to b_i , or if the bid was submitted by an existing bidder.

VI. Efficiency

In this section, we provide a comparison of Wang and Leung's CDA scheme to several implementations of our CDA approach based on group signatures. Table VI shows the amount of work performed during each major stage of the CDA in terms of the number of modular exponentiations. The schemes compared include: Wang and Leung [21] (WL04), Anteniese *et al.* [1] (ACJT00), Camenisch and Lysyankaya [4] (CL02), and Tsudik and Xu [19] (TX03).

The notation L in Table VI denotes the number of times the cut-and-choose protocol is run in Wang and Leung's scheme. Small values of L increase the chances of cheating, whereas higher values decrease the system efficiency. Wang and Leung suggest L should be at least 20, which requires an average of 60 modular exponentiations for registration. The group signature scheme by Tsudik and Xu is clearly the most efficient in terms of registration.

While Wang and Leung's scheme is more efficient for signing and verification, it suffers from the security drawbacks stated in sections II-A and II-B. Furthermore, Wang and Leung require the Auctioneer to store a public key for each bidder, whereas our scheme only requires one public key for verification (i.e., the group public key \mathcal{Y}).

Revocation is measured as the work that must be performed in terms of the number of revoked bidders, or the number of remaining bidders ℓ . It is unclear how Wang and Leung's scheme deals with revoked bidders. Ateniese, Song and Tsudik [2] revoke bidders in the ACJT group signature scheme using a Certificate Revocation List. The Auctioneer must check each bid against the list, which is linear in size of the number of revoked bidders. Alternatively, Camenisch and Lysyankaya, and Tsudik and Xu use a method based on dynamic accumulators where the cost of checking if a bidder has been revoked is constant. With this approach the burden of checking for revocation rests with the bidder.

A. Other Approaches

Instead of using a group signature scheme, we compare some alternate approaches to preventing linkability of bids.

The first approach is to use a *one-time pseudonym*. This means that once a bidder has submitted a bid, he/she must re-register to obtain a new pseudonym. Continuously running the registration protocol of Wang and Leung is more inefficient than using a group signature scheme. It would make the procrastinating attack harder, but would not eliminate it if a bidder were to bid immediately.

Another approach is to issue *multiple pseudonyms* during registration. This prevents a bidder from being profiled, and eliminates the need to run the registration protocol for each bid submitted. However, this method requires the bidder to store a large number of pseudonyms, whereas the group signature approach only requires a single size certificate. Furthermore, this complicates the revocation as now the Auctioneer has to check a large number of pseudonyms for a single bidder.

VII. Conclusion

In this paper we proposed a new scheme for conducting secure and anonymous CDAs. Such a scheme is vital for protecting the security and anonymity of participants who engage in share trading via online brokerage firms. We showed that the only existing CDA proposal by Wang and Leung is inadequate for the task. First of all, its security can be readily subverted by a procrastinating attack. Secondly, their scheme allows bids to be linked and profiles can be created about individual bidders. Finally it is inefficient in the following three areas: registration, bid verification, and bidder revocation. 1) Registration is inefficient due to the reliance on the cut-and-choose protocol; 2) Bid verification requires the Auctioneer to hold a large number of keys (i.e., proportional to the number of bidders); 3) Once a bidder has been traced, it is unclear how the bidder can be revoked.

In direct contrast, our scheme is more flexible and solves all of these problems by utilising any existing secure group signature scheme. The procrastinating attack is ineffective due to the properties of the group signature scheme. Bids cannot be linked to a particular bidder as every signature is different each time a bid is signed, therefore bidders cannot be profiled. While our scheme is slower for bidding and verification, it is efficient in the following respects: 1) Registration does not require expensive cut-and-choose operations; 2) The Auctioneer only requires one key for verifying bids; 3) The cost of revocation can be reduced to a constant. Furthermore, we provide a more comprehensive set of security requirements which are provably secure. The efficiency and security of our CDA scheme rests with the underlying group signature scheme used, and our approach offers the client flexibility in choosing from any group signature scheme.

Clearly much work is required on improving the efficiency of group signature schemes.

References

- [1] Ateniese G., Camenisch J., Joye M. and Tsudik G., "A Practical and Provably Secure Coalition-Resistant Group Signature Scheme," in *Advances in Cryptology - Proceedings of CRYPTO 2000* (M. Bellare, ed.), vol. 1880 of *Lecture Notes in Computer Science*, Springer-Verlag, 2000, 255–270.
- [2] Ateniese G., Song D. and Tsudik G., "Quasi-Efficient Revocation of Group Signatures," in *Proceedings of Financial Cryptography*, (M. Blaze, ed.), vol. 2357 of *Lecture Notes in Computer Science*, Springer-Verlag, 2002, 183–197.
- [3] Boyd C. and Mao W., "Security Issues for Electronic Auctions," Technical Report, Hewlett Packard TR-HPL-2000, 2000.
- [4] Camenisch J. and Lysyanskaya A., "Dynamic Accumulators and Application to Efficient Revocation of Anonymous Credentials," in *Advances in Cryptology - Proceedings of CRYPTO 2002* (M. Yung, ed.), vol. 2442 of *Lecture Notes in Computer Science*, Springer-Verlag, 2002, 61–76.
- [5] Camenisch J. and Stadler M., "Efficient Group Signature Scheme for Large Groups," in *Advances in Cryptology - Proceedings of CRYPTO '97* (S. Burton and J. Kaliski, eds.), vol. 1294 of *Lecture Notes in Computer Science*, Springer-Verlag, 1997, 410–424.
- [6] Chaum D. and van Heyst E., "Group Signatures," in *Advances in Cryptology - Proceedings of EUROCRYPT '91* (D. Davies, ed.), vol. 547 of *Lecture Notes in Computer Science*, Springer-Verlag, 1991, 257–265.
- [7] Franklin M. and Reiter M., "The Design and Implementation of a Secure Auction Service," *IEEE Transactions on Software Engineering*, vol. 22, May 1996, 302–312.
- [8] Friedman D. and Rust J., *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley, 1992.
- [9] Gjrestad S. and Dickhaut J., "Price Formation in Double Auctions," *Games and Economic Behaviour*, vol. 22, 1998, 1–29.
- [10] He M. and Leung H., "An Agent Bidding Strategy Based on Fuzzy Logic in a Continuous Double Auction," in *Proceedings of the fifth International Conference on Autonomous Agents*, 2001, 61–62.
- [11] Jakobsson M. and Juels A., "Mix and Match: Secure Function Evaluation via Ciphertexts," in *Advances in Cryptology - Proceedings of ASIACRYPT 2000* (T. Okamoto, ed.), vol. 1976 of *Lecture Notes in Computer Science*, Springer-Verlag, 2000, 162–177.
- [12] Juels A. and Szydlo M., "A Two-Server, Sealed-Bid Auction Protocol," in *Proceedings of Financial Cryptography*, vol. 2357 of *Lecture Notes in Computer Science*, Springer-Verlag, 2002, 72–86.
- [13] Naor M., Pinkas B., and Sumner R., "Privacy Preserving Auctions and Mechanism Design," in *The 1st Conference on Electronic Commerce*, 1999, 129–139.
- [14] Nguyen K. and Traoré J., "An Online Public Auction Protocol Protecting Bidder Privacy," in *Proceedings of ACISP 2000 – Australasian Conference on Information Security and Privacy* (E. Dawson, A. Clark, and C. Boyd, eds.), vol. 1841 of *Lecture Notes in Computer Science*, Springer-Verlag (Berlin), 2000, 427–442.
- [15] Sandholm T., Blum A., and Zinkevich M., "On-line Algorithms for Market Clearing," in *In Proceedings of the 13th SIAM Symposium on Discrete Algorithms (SODA)*, 2002.
- [16] Tesouro C. and Bredin J., "Strategic Sequential Bidding in Auctions using Dynamic Programming," in *Proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'02)*, 2002, 591–598.
- [17] Trevathan J., "Security, Anonymity and Trust in Electronic Auctions," *Association for Computing Machinery, Crossroads Student Journal*, Spring Edition, vol. 11.3, 2005.
- [18] Trevathan J., Ghodosi H. and Read W., "Design Issues for Electronic Auctions," in *2nd International Conference on E-Business and Telecommunication Networks*, 2005 (to appear).
- [19] Tsudik G. and Xu S., "Accumulating Composites and Improved Group Signing," in *Advances in Cryptology - Proceedings of ASIACRYPT 2003* (C. Lai ed.), vol. 2894 of *Lecture Notes in Computer Science*, Springer-Verlag, 2003, 269–286.
- [20] Viswanathan K., Boyd C., and Dawson E., "A Three Phased Schema for Sealed Bid Auction System Design," in *Proceedings of ACISP 2000 – Australasian Conference on Information Security and Privacy* (E. Dawson, A. Clark, and C. Boyd, eds.), vol. 1841 of *Lecture Notes in Computer Science*, Springer-Verlag (Berlin), 2000, 412–426.
- [21] Wang C. and Leung H., "Anonymity and Security in Continuous Double Auctions for Internet Retail Market," in *the 37th Hawaii International Conference on Systems Sciences*, 2004.
- [22] Wellman M., Wurman P., and Walsh W., "The Michigan Internet AuctionBot: A Configurable Auction Server for Human and Software Agents," in *Second International Conference on Autonomous Agents (AGENTS)*, May 1998, 301–308.
- [23] Wellman M., Wurman P., and Walsh W., "A Parametrization of the Auction Design Space," in *Games and Economic Behavior*, 35(1-2), 2001, 304–338.

Chapter 8

Online Auction Software 2

This chapter concludes the software design of the online auction server. Two papers are presented. The first paper describes the software characteristics of a CDA (see Section 8.1). The second paper contrasts several approaches for clearing a CDA (see Section 8.2) that were investigated as part of the auction server's software design.

8.1 A Software Architecture for Continuous Double Auctions

Implementing an online CDA in software is more complicated than a regular online auction (such as eBay). This paper describes our experiences with implementing an online CDA. The model is presented as an abstraction of the online share trading process, and discusses implementation specific details. The major software components are described, along with web site navigation and object-oriented software design. An online CDA database schema is presented along with a discussion regarding timing issues. The paper shows how bids are cleared and contrast differing matching strategies. Furthermore, it investigates CDA software bidding agents, presents an agent application programming interface, and gives a description of different bidding strategies.

“A Software Architecture for Continuous Double Auctions” [15] is published in the Proceedings of the International Association for Development of the Information Society (IADIS) – International Conference on Applied Computing 2007. IADIS is peer-reviewed by an international panel of experts and was held in Salamanca, Spain. IADIS has the following goals:

1. To develop the cooperation and solidarity between its associates, by developing initiatives in the Information Society domain;
2. To promote the study, research and dissemination of news related to the Information Society;
3. To make available to its associates information and bibliography about the Information Society;
4. To organize working groups, to research, study, develop and analyze issues related to the Information Society;
5. To publish magazines, journals or other documents of significant interest;
6. To organize meetings, seminars and conferences;

7. To promote people training with the goal of integrating them in the Information Society;
8. To promote the exchange and cooperation with national and foreign associations and entities, which seek the same goals.

THIS PAPER HAS BEEN REMOVED DUE TO COPYRIGHT RESTRICTIONS

8.2 Variable Quantity Market Clearing Algorithms

Market clearing is the process of matching buy and sell bids in securities markets. The allocative efficiency of such algorithms is important, as the Auctioneer is typically paid a commission on the number of bids matched and the volume of quantity traded. Previous algorithms have concentrated on price issues. This paper presents several market clearing algorithms that focus solely on allocating quantity among matching buy and sell bids. The goal is to maximise the number of bids matched, while at the same time minimise the amount of unmatched quantity. The algorithms attempt to avoid situations resulting in unmarketable quantities (i.e., quantities too small to sell). Algorithmic performance is tested using simulated data designed to emulate the Australian Stock Exchange (ASX) and other world stock markets. Results show that it is difficult to avoid partial matchings as the complexity of doing so is NP-complete. The optimal offline algorithm for partial quantity matching is used as a benchmark to compare online matching strategies. Three algorithms are presented that outperform the ASX's strategy by increasing the number of bids matched, the amount of quantity matched, and the number of bids fully matched.

“Variable Quantity Market Clearing Algorithms” [8] is published in the proceedings of the International Conference on e-Business (ICE-B 2006). ICE-B 2006 was held in Setúbal, Portugal and was organised by INSTICC. Sponsors included IBM, Setúbal Polytechnic Institute, IEEE Systems, Man and Cybernetics (SMC) Society, and ACM SIGMIS. The major goal of this conference was to bring together researchers and developers from academia and industry working in areas related to e-business to stimulate a debate with a special focus on telecommunications networks. The conference received 326 submissions in total, with contributions from 53 different countries. To evaluate each submission, a double blind paper evaluation method was used: each paper was reviewed by at least two internationally known experts from the Program Committee, and more than 95% of the papers had 3 or more reviews. In the end, 98 papers were selected to be published and presented as full papers. This corresponds to a 30% acceptance ratio. This paper is available online via citeseer ¹ and is referenced by DBLP ².

I was a session chair at ICE-B 2006. This paper was initially conceived in collaboration with Associate Professor Bruce Litow. Associate Professor Litow's research strength is theoretical Computer Science. Publication was made possible by funding from the School of Mathematical and Physical Sciences.

“Variable Quantity Market Clearing Algorithms” [19] has also been selected to appear in the “ICETE 2006 Best Papers” book published by Springer Verlag. ICE-B is part of the International Conference on E-Business and Telecommunications Engineers (ICETE). Out of four independent conferences held under the ICETE banner, only 28 of the best papers were accepted for publication in the book. This corresponds to an 8% acceptance rate.

¹<http://citeseer.ist.psu.edu/>

²<http://www.informatik.uni-trier.de/~ley/db/>

VARIABLE QUANTITY MARKET CLEARING ALGORITHMS

Jarrold Trevathan

School of Mathematical and Physical Sciences

James Cook University

Email: jarrod.trevathan@jcu.edu.au

Wayne Read

School of Mathematical and Physical Sciences

James Cook University

Email: wayne.read@jcu.edu.au

Keywords: Continuous double auctions, online auctions, share trading, competitive analysis.

Abstract: Market clearing is the process of matching buy and sell bids in securities markets. The allocative efficiency of such algorithms is important, as the Auctioneer is typically paid a commission on the number of bids matched and the volume of quantity traded. Previous algorithms have concentrated on price issues. This paper presents several market clearing algorithms that focus solely on allocating quantity among matching buy and sell bids. The goal is to maximise the number of bids matched, while at the same time minimise the amount of unmatched quantity. The algorithms attempt to avoid situations resulting in unmarketable quantities (i.e., quantities too small to sell). Algorithmic performance is tested using simulated data designed to emulate the Australian Stock Exchange (ASX) and other world stock markets. Our results show that it is difficult to avoid partial matchings as the complexity of doing so is NP-complete. The optimal offline algorithm for partial quantity matching is used as a benchmark to compare online matching strategies. We present three algorithms that outperform the ASX's strategy by increasing the number of bids matched, the amount of quantity matched, and the number of bids fully matched.

1 INTRODUCTION

Securities markets such as the New York Stock Exchange¹ and the Australian Stock Exchange² (ASX), employ a form of auction referred to as a *Continuous Double Auction* (CDA). A CDA has many buyers and sellers continuously trading a commodity. Buy and sell bids accumulate over time and must be cleared. The method by which buy and sell bids are matched is referred to as a *market clearing algorithm*. In general, two bids can only be matched if: 1) both bids are currently active (i.e., they haven't expired or previously been cleared); and 2) the price of a buy bid equals or exceeds the price of a sell offer.

The efficiency and performance of a clearing algorithm is important. An algorithm must be able to cope with large numbers of bids, and make timely decisions which maximise the benefits for buyers and sellers. Furthermore, the Auctioneer/Broker typically gains commission on the *number of bids cleared*, and the *volume of quantity traded*. As a result, the algorithm must also strive to maximise both of these fac-

tors.

Stock exchanges have been fully automated since the early 1990s (see (Economides and Schwartz, 1995)). The ASX uses a computerised clearing system referred to as the Stock Exchange Automated Trading System (SEATS). SEATS imposes a strict time-based priority on matching bids. Bids are ordered according to price, and are then matched based on their arrival times. Larger bids are not given priority over small bids.

Alternate strategies for market clearing have been discussed by (Wellman and Wurman, 1999; Sandholm and Suri, 2001). (Sandholm *et al.*, 2002) show that in some situations, the Auctioneer can increase the profit from a sale (i.e., the price difference between a buy and sell bid). This is achieved by not matching bids immediately, but rather waiting for a better match to possibly eventuate. (Sandholm *et al.*, 2002) also describe how profit producing matches can subsidise loss producing matches to increase the total number of bids matched.

The market clearing model used by (Sandholm *et al.*, 2002) mainly attempts to maximise the amount of surplus generated by the matching process. In doing so, the model only considers price, and assumes that

¹<http://www.nyse.com>

²<http://www.asx.com.au>

the quantity of each bid is *one* unit. If a bidder desires to bid for more than one unit, then they must enter a number of bids equal to the amount of the quantity. (e.g., five separate bids are required to bid for a quantity of five units.)

This approach is not very practical when the amount of quantity transacted is large. For example, in the case where a bidder desires 10,000 units, it is unlikely they would be willing to expend time and effort to submit 10,000 bids. While software bidding agents and other automated methods could be used to alleviate this situation, there are further issues regarding allocating quantity among bids. For example, a bidder may desire n units, but the market is only able to clear ρ units, where $\rho < n$. This may leave the bidder with an *unmarketable quantity*. An unmarketable quantity, is a quantity that is too small to sell, after taking into account Auctioneer commission, and other associated costs.

In this paper, we propose the idea of a *variable quantity market clearing algorithm*. Once bids have been ordered according to price, a variable quantity market clearing algorithm is used to efficiently allocate quantity among matching buy and sell bids. The algorithm attempts to match up as many bids as possible, with as little or no unmatched quantity outstanding. The primary goal is to avoid situations that result in unmarketable quantities.

This paper presents several variable quantity market clearing algorithms. The first algorithm shows why it is difficult in practice to avoid unmarketable quantities. The second algorithm gives the optimal offline solution in terms of avoiding unmarketable quantities. The third algorithm is online, and is the approach used by SEATS. The remaining algorithms are online, and try to outperform Algorithm 3, using differing strategies including: waiting until a bid is ready to expire before matching, subsidising short falls in allocation, and giving priority to bids with smaller quantities.

The algorithms have been tested on simulated data designed to emulate the workings of the ASX. Each algorithm is assessed according to the number of bids matched, the volume of quantity traded and how much unmarketable quantity is produced. We show that it is possible to out-perform SEATS in terms of these goals.

This paper is organised as follows: The CDA model and goals of the algorithms are discussed in Section 2. Section 3 presents several market clearing algorithms for matching variable quantities of an item. A comparison of the algorithms is given in Section 4, and Section 5 provides some concluding remarks.

2 PRELIMINARIES

This section presents a CDA model for describing variable quantity market clearing algorithms. The goals for a clearing algorithm are discussed, and basic statistics are introduced for measuring how an algorithm performs in terms of these goals.

2.1 Model

The algorithms presented in this paper are based on a *temporal clearing model*. This consists of a set of buy bids, B , and a set of sell bids, S . Each bid $v \in B \cup S$ has the components $(type, t_i, t_j, p, q)$.

$type = \{buy, sell\}$ denotes the type of the bid. It is common in securities markets to refer to a buy bid as a *bid* and a sell bid as an *offer*.

A bid, v , is introduced at time $t_i(v)$, and removed at time $t_j(v)$, where $t_i(v) < t_j(v)$. A bid, v , is said to be alive in the interval $[t_i(v), t_j(v)]$. To be a candidate for a particular matching, two bids must have a period where their arrival and expiration times overlap. Two bids $v, v' \in B \cup S$ are said to be *concurrent*, if there is some time when both are alive simultaneously.

p denotes the price of a bid. In order to run the clearing algorithm, bids are first ordered according to price. The definition of concurrency now extends to two bids that met the criteria for matching based on price. This allows us to concentrate on matching quantities rather than prices. The problem now becomes the opposite extreme of the price-matching problem from (Sandholm *et al.*, 2002).

$q \in [q_{min}, q_{max}]$ denotes the quantity, and $q(v)$ is the quantity desired by bid v . q must be greater than zero and an integer, i.e., it is not possible to buy or sell a fraction of a quantity.

The temporal bidding model is abstracted as an incomplete interval graph. An incomplete interval graph is a graph $G = (V, E)$, together with two functions t_i and t_f from V to $[0, \infty]$ such that:

1. For all $v \in V$, $t_i(v) < t_f(v)$. (i.e., the entry time is less than the exit time.)
2. If $(v, v') \in E$, then $t_i(v) \leq t_f(v')$ and $t_i(v') \leq t_f(v)$. (i.e., bids are concurrent.)

An incomplete interval graph can be thought of as an abstraction of the temporal bidding problem. The fact that bids come in two types (buy and sell) and have prices attached to them is ignored. Instead a black box “E” is used, that given two bids v, v' , outputs whether or not they can be matched. This generalisation provides a simple framework for describing and developing clearing algorithms.

2.2 Goals

The quantity matching algorithms presented in this paper have the following goals:

Maximise Number of Matches

The first goal is to maximise the number of matches between buy and sell bids. This is important as the Auctioneer typically gains a commission based on the number of bids matched.

The *Bid Match Ratio* (BMR) is used to measure the number of bids matched, α , in relation to the total number of bids, n . This is calculated as:

$$BMR = \alpha/n \times 100$$

where $0 \leq BMR \leq 100$.

Maximise Volume of Quantity Matched

The second goal is to maximise the amount of quantity matched. The Auctioneer also typically gains a commission proportional to the volume of quantity cleared.

The *Quantity Match Ratio* (QMR) is used to measure the amount of matched quantity, δ , in relation to the total quantity, γ . This is calculated as:

$$QMR = \delta/\gamma \times 100$$

where $0 \leq QMR \leq 100$.

Maximise Full Quantity Matches

A *full* match occurs when a bid has had its entire quantity matched and cleared. A *partial* match occurs when a bid which is in the process of being matched, expires with an outstanding quantity that hasn't been filled. The Auctioneer must strive to satisfy the entire quantity of a bid, so that a bidder is not left with an unmarketable quantity.

The *Full Match Ratio* (FMR) examines the bids that were fully matched, ϵ , against the total number of matches, ζ . This is calculated as:

$$FMR = \epsilon/\zeta \times 100$$

where $0 \leq FMR \leq 100$.

2.3 Analysing Efficiency

Within both Computer Science and Finance, many problems reduce to trying to predict the future. For example, cache/virtual memory management, process scheduling, or predicting future returns for an asset. Such problems become trivial if the future is known (i.e., the stream of future memory requests or tomorrow's newspaper), but typically we only have access to the past.

An offline problem provides access to all the relevant information to compute a result. An online problem continually produces new input and requires answers in response. Offline problems have the benefit

of perfect knowledge, and such they generally outperform online problems (if designed properly).

An *offline* clearing algorithm learns of all bids up front. That is, all bids must be submitted before a closing time. The algorithm is then able to match bids at its discretion. An *online* clearing algorithm only learns about bids as they are introduced over time. The online algorithm has the added complexity of bids expiring before they can be matched.

Securities markets employ both types of algorithms. For example, the online algorithm is used during trading hours and the offline algorithm is used after hours while bids accumulate over night.

Competitive analysis allows an online algorithm to be compared based on its ability to successfully predict the future. The efficiency of an online solution is compared to the optimal offline solution. The closer an online algorithm performs to the optimal offline algorithm, the more 'competitive' the algorithm is.

An algorithm, A , is said to be *c-competitive* if there exists some constant b , such that for every sequence of inputs σ :

$$cost_A(\sigma) \leq c \cdot cost_{OPT}(\sigma) + b$$

where OPT is the optimal offline algorithm. In developing an online algorithm, the goal is to attain a *competitive ratio*, c , as close to one as possible. The worse the performance of an algorithm, the larger c is.

In this paper, an optimal offline solution is presented for clearing variable quantities. Several online strategies are discussed, and their performance is compared based on their competitive ratios. Related literature on how competitive analysis has been applied to online auctioning can be found in (El-Yaniv *et al.*, 1992; Lavi and Nisan, 2000; Bagchi *et al.*, 2001).

3 QUANTITY CLEARING ALGORITHMS

This section presents several market clearing algorithms for matching variable quantities of an item.

3.1 Algorithm 1

The initial goal for this algorithm is to either match quantities entirely, or not at all (i.e., bids are *indivisible*). This effectively eliminates the possibility of unmarketable quantities.

For example, a buy and sell bid each for 1 unit, can be matched. In addition, a buy bid for 2 units can be matched to two sell bids that are for 1 unit each. However, if a buyer is demanding 2 units, and there is a seller supplying only 1 unit, then neither of the bids can be matched.

Matching indivisible quantities is similar to the Knapsack problem. Consider the case where there is a buy bid v for a quantity $q(v) = n$. In order to match this bid, the algorithm is required to select either a single sell bid v' offering n units, or some combination of n or less sell bids, where the collective quantity on offer sums exactly to n . The complexity of the algorithm for matching indivisible quantities is dominated by this step.

The problem of trying to find a subset of integers from a given set, that sums to a particular value, is referred to as the *subset sum problem*. The subset sum problem is considered to be hard, and there are no known algorithms to efficiently solve it. The quandary of matching indivisible quantities can be reduced to the subset sum problem. The subset sum problem is NP-complete, and therefore the indivisible quantity matching problem is also NP-complete. As a result, it is not feasible to construct an efficient algorithm that does not deliver unmarketable quantities.

Even if this algorithm were practical, it does not necessarily perform well in terms of the number of bids matched, as it is too restrictive. The costs of unmarketable quantities must be weighed against the benefits of relaxing the indivisibility constraint. Doing so allows the clearing process to benefit the majority of bidders, while at times delivering an undesirable result to a minority. The problem now becomes how to limit the extent of unmarketable quantities.

3.2 Algorithm 2

This algorithm is offline and allows bids to be *divisible*. A particular bid is matched with as many other candidate bids as required to satisfy it. If there is not enough available quantity, the bid is considered as *partially matched*. A partial matching can result in bidders holding unmarketable quantities. The goal of this algorithm is to match as many bids as possible and minimise partial matchings.

A greedy strategy is employed which successively subtracts the smaller bid quantity from the larger opposite bid quantity. The algorithm keeps track of the current unmatched buy and sell quantities at each stage of the algorithm using two variables, α_b and α_s . Once a particular bid has been allocated its exact quantity, it is cleared (i.e., moved to the set M). The algorithm is as follows:

1. $\alpha_b = \alpha_s = 0$.
2. While there are more vertices in G
 - (a) if α_b and $\alpha_s = 0$ then
 - i. get next v and v' from G
 - ii. if $q(v) > q(v')$ then
 - $\alpha_b = q(v)$
 - iii. else
 - $\alpha_s = q(v')$

- (b) else if $\alpha_b > 0$ then
 - i. get next v' from G
 - ii. if $\alpha_b > q(v')$ then
 - A. $\alpha_b = \alpha_b - q(v')$
 - B. place edge between v and v' , move v' to M
 - iii. else
 - A. $\alpha_s = q(v') - \alpha_b$
 - B. place edge between v and v' , move v to M
- (c) else if $\alpha_s > 0$ then
 - i. get next v from G
 - ii. if $\alpha_s > q(v)$ then
 - A. $\alpha_s = \alpha_s - q(v)$
 - B. place edge between v and v' , move v to M
 - iii. else
 - A. $\alpha_b = q(v) - \alpha_s$
 - B. place edge between v and v' , move v' to M

As all bids are concurrent, the proposed solution is equivalent to summing the volumes of buy and sell bids, and subtracting the smaller from the larger. This algorithm is the optimal solution for matching variable quantities, and is the basis for the operation of the forthcoming online algorithms.

3.3 Algorithm 3

This algorithm is online and uses the same strategy as Algorithm 2. Bids have entry and expiration times. When a bid is introduced, it is matched with as many other bids as possible. However, when expiration time is reached, the bid is cleared regardless of whether it has been fully matched. That is, if a bid is in the process of being matched when it expires, its outstanding quantity remains unfilled.

This is the actual approach used by SEATS and most of the world's securities markets. It is simple, fair and performs relatively well. However, this algorithm performs significantly worse than the previous algorithm, and can result in many bids expiring partially matched.

3.4 Algorithm 4

Algorithm 4 aspires to out-perform the previous algorithm. When a bid expires in Algorithm 3, there may be a significant amount of residual unmatched quantity. In addition, bids that arrive later, but expire earlier have to wait on earlier bids, with later expiration times. Algorithm 4 modifies the previous algorithm by waiting till a bid is about to expire, before matching it with as many other bids as possible based on expiry time.

3.5 Algorithm 5

Algorithm 5 uses an inventory of quantity to offset the unfilled portions of partially matched bids. We refer to this as *subsidising*.

To acquire an inventory, the Auctioneer must first collect surplus quantity. A *surplus* occurs when the quantity on offer is greater than the quantity being bid for. If a buy bid demands a quantity less than the quantity offered by a sell bid, then the Auctioneer pays for the surplus. The surplus quantity is placed in an inventory that can be used at a later date to subsidise shortfalls in allocation. A *shortage* occurs when there is less quantity on offer than the amount bid for.

Determining the extent to subsidise shortages requires choosing a threshold, which is the maximum amount that can be subsidised. This is denoted by θ .

The algorithm proceeds in a similar manner to the previous online algorithms. However, when a bid is about to expire, it becomes a candidate for subsidisation. Let I denote the current inventory of quantity held by the algorithm. The subsidisation process is as follows:

1. if v is type *sell* then
 - (a) if $(I + q(v) < \theta)$ then
 $I = I + q(v)$
 - (b) else
 $temp = \theta - I$
 $I = I + temp$
 $q(v) = q(v) - temp$
2. else if v is type *buy* then
 - (a) if $(I < q(v))$ then
 $q(v) = q(v) - I$
 - (b) else
 $I = I - q(v)$
 $q(v) = 0$

3. Move v to M

The choice of θ depends on the risk the Auctioneer is willing to take. If θ is set too small, the Auctioneer will not be able to significantly influence the clearing process. However, if θ is set too large, the Auctioneer might be left holding a large quantity at the close of trade, which is undesirable.

The *Remaining Inventory Ratio* (RIR) is used to measure the extent of remaining quantity held by the Auctioneer at the close of trade. The RIR is calculated as follows:

$$RIR = I/\theta \times 100$$

where $0 \leq RIR \leq 100$.

3.6 Algorithm 6

A problem with the greedy approach of Algorithms 2 and 3 is that a large number of smaller bids may be

waiting on a earlier, larger bid to clear.

In economic systems it is usually the case that a smaller number of individuals own the most. Likewise, in share trading there tend to be more bids for smaller quantities compared to larger bids. Large bids are often due to financial institutions such as superannuation schemes or managed funds that pool the capital of many smaller investors. An Auctioneer can take advantage of the above situation by clearing the smaller bids first. This will increase the number of bids matched while leaving the volume traded unchanged.

Algorithm 6 gives priority to smaller bids. That is, if a large bid is in the process of being matched, it is 'pre-empted' when a smaller bid arrives. This is analogous to the problem of process scheduling where many processes compete for CPU time. In shortest job first scheduling, the process with the shortest time is given priority to use the CPU.

In terms of market clearing, quantity represents CPU time and bidders are the processes. However, matching bids is two sided and therefore more complicated than process scheduling. That is, the set of buy bids represents a set of processes, and the set of sell bids represents another set of processes. When a large buy bid is matched to several smaller sell bids, the buy bid is equivalent to the CPU for that instance in time (and vice versa).

4 COMPARISON

This section provides a comparison of the market clearing algorithms. Each algorithm is assessed on its ability to achieve the goals outlined in Section 2.2. These goals are: maximise the number of matched bids (BMR), maximise the volume of quantity matched (QMR), and maximise the number of fully matched bids (i.e., avoid partial matchings)(FMR).

The *Research Auction Server* at James Cook University, is an online auction server used to conduct research into online auctioning [see (Trevathan and Read, 2006)]. The auction server contains a simulation environment for testing computational aspects related to auctioning online. This includes emulating the workings of the ASX by generating the kind of bidding data that exists in share markets. The clearing algorithms were tested in this setting using the simulated data.

The input parameters for a test are the number of bids, n , the maximum quantity allowable for a bid, q_{max} , and the total time, t . Time is split into discrete units representing seconds. Entry and Exit times are randomly generated between time period one and t . Bids are randomly allocated quantities between one




Figure 1: Algorithmic Performance on Simulated Data

Table 1: Comparison of Variable Quantity Market Clearing Algorithms

	Type	BMR	QMR	FMR	RIR	Competitive Ratio
Alg. 2	offline	$0.43 \ln(n)$	$0.47 \ln(n)$	$0.05 \ln(n)$	-	Optimal
Alg. 3	online	$1.05 \ln(n)$	$1.23 \ln(n)$	$0.29 \ln(n)$	-	2.44
Alg. 4	online	$0.51 \ln(n)$	$0.71 \ln(n)$	$0.30 \ln(n)$	-	1.18
Alg. 5a	online	$0.81 \ln(n)$	$0.91 \ln(n)$	$0.21 \ln(n)$	50.61%	1.88
Alg. 5b	online	$0.43 \ln(n)$	$0.48 \ln(n)$	$0.11 \ln(n)$	50.13%	1.00
Alg. 6	online	$0.77 \ln(n)$	$1.20 \ln(n)$	$0.29 \ln(n)$	-	1.79

and q_{max} . For now, it is assumed that entry/exit times and quantities are uniformly distributed among bidders, and there is an even number of buy and sell bids (i.e., $\# \text{ buy} = \# \text{ sell} = n/2$).

Figure 1 shows the individual results for each algorithm (except Algorithm 1). Each algorithm was tested using varying numbers of bids up to a maximum of $n = 10,000$, with $q_{max} = 1000$. The online algorithms are shown for a six-hour time period (i.e., $t = 21,600$ seconds). This is consistent with the trading hours of the ASX and most share markets.

In general, all algorithms run in a time proportional to n . Furthermore, the larger the value of n , the better the performance. Additionally, online algorithms tended to perform better with smaller values of t . The size of q_{max} does not significantly affect the running time or the performance of the algorithms.

For each algorithm, trendlines are calculated using the *method of least squares*. The resulting equations for each algorithm are listed in Table 1. The smaller the value of the coefficient of $\ln(n)$ for the BMR, QMR and FMR, the better the performance. Competitive analysis is used to compare the performance of the online algorithms to the optimal offline algorithm. The BMR is used to determine an algorithm's competitive ratio.

Algorithm 2 is the optimal offline strategy for matching divisible bids. Only a small percentage of bids were partially matched. This algorithm is the benchmark to which all online algorithms are compared. This algorithm is $1 - \text{competitive}$.

Algorithm 3 is online, and orders bids strictly according to entry time (the approach used by SEATS). This algorithm performs significantly worse than the offline algorithm. This algorithm is $2.44 - \text{competitive}$.

Algorithm 4 waits until a bid is about to expire before matching it. This strategy is a significant improvement on Algorithm 3. The reason for this is that by waiting, the algorithm is essentially acting similar to the offline algorithm. Although this algorithm doesn't have perfect knowledge about future bids, delaying matching allows the algorithm to gather more information, which it can use to optimise its matching decision. However, Algorithm 4's FMR does not fare much better than Algorithm 3. This algorithm is $1.18 - \text{competitive}$.

Algorithm 5 uses an inventory of quantity to subsidise deficit quantity trades. Two tests were conducted with differing values of θ . The first test (referred to as Algorithm 5a), examined the effect of minimal subsidisation. The second test (referred to as Algorithm 5b), used excessive subsidisation. Both tests were also assessed on the amount of inventory remaining at the end (i.e., the RIR, see Section 3.4).

Algorithm 5a uses minimal subsidisation where $\theta = 1000$. This algorithm improved upon Algo-

rithm 3 (i.e., SEATS) with regard to its BMR and QMR. This algorithm achieved a better FMR than both previous online algorithms. This algorithm is $1.88 - \text{competitive}$.

Algorithm 5b uses excessive subsidisation where $\theta = 5000$ ($5 \times q_{max}$). This algorithm's performance approaches Algorithm 2 in terms of BMR and QMR. It also attains an excellent FMR. This algorithm is $1 - \text{competitive}$. If subsidisation were without bound, eventually all bids would be matched.

The amount of remaining inventory (i.e., RIR) for Algorithm 5a and 5b were 50.61% and 50.13% respectively. This shows that over time, the clearing algorithm will always be holding an inventory that is half full (i.e., $\theta/2$). While excessive subsidisation may achieve significant results, the practicality of subsidising must be weighed against the level of risk the Auctioneer is willing to take.

Algorithm 6 prioritises bids with smaller amounts of quantity. This algorithm outperforms Algorithm 3 (SEATS) in terms of its BMR, and improves on using minimal subsidisation. The QMR and FMR are marginal improvements on Algorithm 3. This algorithm is $1.79 - \text{competitive}$. This result is consistent with the goals of the algorithm. That is, we strived for an increase in the BMR by matching a larger number of smaller bids. In doing so, the amount of matched quantity would be roughly the same.

5 CONCLUSIONS

A market clearing algorithm's performance greatly affects the revenue earned by the Auctioneer, and the welfare of the bidders. Previous literature on market clearing only addresses price issues, and neglects concerns regarding allocative efficiency.

This paper presents several market clearing algorithms that focus solely on allocating quantity among matching buy and sell bids. The algorithms attempt to avoid situations resulting in unmarketable quantities (i.e., quantities too small to sell).

We show that it is difficult to avoid partial matchings, as the complexity of doing so is NP-complete. The problem of matching bids with indivisible quantities reduces to the subset sum problem. The subset sum problem is a renowned NP-complete problem. As a result, an efficient algorithm cannot be constructed to avoid partial matchings.

An optimal offline algorithm is presented for matching bids with divisible quantities. The algorithm employs a greedy strategy. Each bid is matched with as many other bids as required to satisfy it. This approach achieves a high match rate. However, the algorithm can result in a limited number of bidders receiving partial matchings.

SEATS and most of the world's stock exchanges use an online version of the previous algorithm. Bids are strictly ordered by price and time. Larger bids are not given priority over smaller bids. This algorithm is simple and fair. However, it performs significantly worse than the optimal offline solution.

We propose several alternate methods for clearing variable quantities. The goal is to out-perform the approach used by SEATS. Algorithmic performance is tested using simulated data designed to emulate the ASX. Competitive analysis is used to compare the performance of an online algorithm to the optimal offline solution.

The first of our proposed algorithms showed that there is some benefit in waiting rather than matching with the first available quantity. Waiting until a bid is about to expire before matching, makes the algorithm function more like its offline counterpart. This algorithm performs significantly better than SEATS.

The next algorithm collects surplus quantity to subsidise shortfalls in allocation. With minimal subsidisation, this algorithm can deliver less partial matchings than SEATS. With excessive subsidisation, this algorithm can approach the efficiency of the optimal offline solution. However, the level of subsidisation must be weighed against the potential to be left holding a large inventory at the end of matching.

Alternately, the final algorithm gives priority to bids with smaller quantities. If a bid is in the process of being matched, it is pre-empted by a bid with a smaller quantity. This approach strongly outperforms SEATS in terms of the number of bids matched. However, it only offers a minor improvement in delivering less partial matchings.

In a rigid environment such as a share market, these mechanisms may not be deemed as initially fair. However, our results show that over time, the proposed algorithms can attain a better outcome than SEATS.

The tests assume there are an even number of buy and sell bids with a uniform distribution of quantity. In reality this would not occur. Increasing the number of either type essentially increases the volume on offer for one type in relation to the other. This degrades performance regardless of the algorithm employed. In the extreme case there will be all of one type and none of the other, which in this case the BMR and QMR would be zero. Having an even number of each type of bid is a neutral point. Skewing the number of bids one way or the other is detrimental to performance.

Bids are uniformly distributed across time. In reality, there may be periods of high bidding volume and also low volume periods. Future work involves using a Poisson probability distribution to model the frequency of the bids. This should help show how the algorithms perform on bursty data.

It would be intuitive to test the algorithms on real

stock market data. However, we have found such data difficult to obtain. There exist many commercial securities market data providers such as Bourse Data³ who sell real-time and historical market data. However, the market depth provided only lists the aggregate quantity at a price level and not the individual bids.

REFERENCES

- Bagchi, A., Chaudhary, A., Garg, R., Goodrich, M. and Kumar, V. (2001). Seller-focused Algorithms for Online Auctioning, in *Proceedings of WADS 2001*, 135-147.
- Economides, N. and Schwartz, R. (1995). Electronic Call Market Trading, *Journal of Portfolio Management*, vol. 21, no. 3, 10-18.
- Kao, M. and Tate, S. (1999). Online Difference Maximisation, *SIAM Journal of Discrete Mathematics*, vol. 12, no. 1, 78-90.
- El-Yaniv, R., Fiat, A., Karp, R. and Turpin, G. (1992). Competitive Analysis of Financial Games, in *Proceedings of the 33rd Symposium on Foundations in Computer Science*, 327-333.
- Lavi, R. and Nisan, N. (2000). Competitive Analysis of Incentive Compatible Online Auctions, in *Proceedings of the 2nd ACM Conference on Electronic Commerce*, 233-241.
- Sandholm, T. and Suri, S. (2001). Market Clearability, in *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, 1145-1151.
- Sandholm, T., Blum, A. and Zinkevich, M. (2002). Online Algorithms for Market Clearing, in *Proceedings of the 13th SIAM Symposium on Discrete Algorithms (SODA)*, 971-980.
- Trevathan, J. and Read, W. (2006). RAS: a system for supporting research in online auctions, *ACM Crossroads*, ed. 12.4, 23-30.
- Wellman, M. and Wurman, P. (1999). A Parameterization of the Auction Design Space, in *Proceedings of the Second International Conference on Autonomous Agents (AGENTS)*, 301-308.

³<http://www.boursedata.com.au>

Chapter 9

Conclusions

This thesis examines privacy and security concerns in online auctions. Existing literature tends to focus on sealed bid auctions, whereas our research primarily deals with online English auctions and CDAs. Major types of undesirable and fraudulent behaviour are discussed. An online auction server was constructed to aid in developing and evaluating security and anonymity mechanisms (i.e., RAS). The auction server's software components are documented, and test results are given.

An online English auction security model is presented. This model is more comprehensive than any existing proposal. Present security methods employed by commercial online auctioneers are insufficient. Such methods often rely on insurance and Escrow techniques, which also can be prone to security scams. The proposed model uses cryptographic mechanisms (i.e., a group signature scheme) to ensure bid authenticity and to prevent a bidder from repudiating having made a bid. Furthermore, the scheme is publicly verifiable, which means that all parties can be confirmed as having followed the auction protocol correctly. The scheme is also flexible in that its efficiency can be improved with advances in the underlying cryptographic techniques. An individual can bid anonymously, unless they misbehave, in which case two independent companies can recover his/her identity. The scheme is able to ensure correct auction timing, and prevents the Auctioneer from blocking bids.

Another major problem in online auctions is shill bidding, where fake bids are used to artificially inflate the price of an auction. A method to detect shill bidding in online auctions has been developed. Bidders are issued with a score based on their bidding behaviour. This is referred to as a shill score. A bidder's shill score indicates the likelihood that s/he is behaving in shilling. Colluding shill bidders can engage in tactics to reduce the chances of being caught. These strategies have been identified, and a method devised to highlight which bidders are most likely to be in collusion. This is referred to as the collusion graph. The final stage in this part of the research is to assign sellers a global rating, which indicates whether a seller exhibits price inflating behaviour. This will also indicate if they are part of a colluding ring of bidders/sellers. All proposed methods have been successfully tested in both simulated and real auctions conducted on RAS. The shill detection methods have been tested on commercial auction data, which has uncovered some extremely suspicious behaviour.

In addition, automated bidding agent security concerns are discussed. Software bidding agents that follow shill bidding strategies are presented. The malicious agents were constructed to aid in developing the proposed shill detection techniques. The simple shill agent incrementally increases an auction's price, forcing legitimate bidders to submit higher bids in order to win an item. The agent ceases bidding when the desired profit from shilling has been attained, or in the case that it is too risky to continue bidding without winning the auction. The adaptive shill agent is used over a series of auctions with substitutable

items, and can revise its strategy based on bidding behaviour in past auctions. The agent uses the EC algorithm, which is a novel approach for determining the minima or maxima in a noisy dataset. The ability of the agent to inflate the price has been tested in a simulated marketplace and experimental results are presented. The EC algorithm's superiority over other extrema avoidance algorithms is shown.

Privacy and security for CDAs was discussed, and how this impacts online share trading. CDAs have more complicated rules than English auctions. Furthermore, the additional parties in the online share trading process makes it difficult to define clear security and privacy policies. Regulatory authorities can not be relied upon to police all e-commerce crime related to online share trading. A model for conducting anonymous and secure CDAs is presented. This can be applied to online share trading. Similar to the English auction security model, the proposed CDA scheme uses a group signature. This guarantees it the same set of privacy and security characteristics. This scheme is the first of its kind to be applied to share markets.

Several new methods for clearing CDAs are presented, which are more efficient than the existing methods used in financial markets. These market clearing algorithms focus solely on allocating quantity among matching buy and sell bids. The goal is to maximise the number of bids matched, while at the same time minimise the amount of unmatched quantity. The algorithms attempt to avoid situations resulting in unmarketable quantities (i.e., quantities too small to sell). Algorithmic performance is tested using simulated data designed to emulate the Australian Stock Exchange (ASX) and other world stock markets. Results show that it is difficult to avoid partial matchings as the complexity of doing so is NP-complete. The optimal offline algorithm for partial quantity matching is used as a benchmark to compare online matching strategies. The proposed algorithms outperform the ASX's strategy by increasing the number of bids matched, the amount of quantity matched, and the number of bids fully matched.

Future work involves devising a more efficient group signature scheme. This would allow both the English and CDA security models to perform signing, verification and identity revocation quicker. Other future work involves testing the skill score on a large scale involving millions of commercial auctions. It is hoped that this will further support claims regarding the practicality of the proposed skill detection techniques. There are also plans to create several more malicious bidding agents that have applications in CDAs and that attack security protocols. It is envisaged that the EC algorithm could possibly be used in conjunction with a neural network, to provide a more accurate measure of share price prediction. Furthermore, it would be intuitive to implement and trial the CDA security model, and clearing algorithms with a commercial online broker, banks/financial institutions and the stock exchanges.

Bibliography

- [1] Trevathan, J., Electronic Auctions - Literature Review, *James Cook University*, 2004.
- [2] Trevathan, J., Security, Anonymity and Trust in Electronic Auctions, *ACM Crossroads*, 11.3, 3–9, 2005.
- [3] Trevathan, J. and McCabe, A., Remote Handwritten Signature Authentication, In *Proceedings of the 2nd International Conference on E-Business and Telecommunication Networks (ICETE)*, 335–339, 2005.
- [4] Trevathan, J., Ghodosi, H. and Read, W., Design Issues for Electronic Auctions, In *Proceedings of the 2nd International Conference on E-Business and Telecommunication Networks (ICETE)*, 340–347, 2005.
- [5] Trevathan, J., Ghodosi, H. and Read, W., An Anonymous and Secure Continuous Double Auction scheme, In *Proceedings of the 39th Hawaii International Conference on System Sciences (HICSS)*, page 125 (1-12), 2006.
- [6] Trevathan, J. and Read, W., RAS: A System for Supporting Research in Online Auctions, *ACM Crossroads*, 12.4, 23–30, 2006.
- [7] Trevathan, J. and Read, W., Secure Online English Auctions, In *Proceedings of the International Conference on Security and Cryptography (SECRYPT)*, pages 387–396, 2006.
- [8] Trevathan, J. and Read, W., Variable Quantity Market Clearing Algorithms, In *Proceedings of the International Conference on e-Business (ICE-B)*, pages 126–133, 2006.
- [9] Trevathan, J. and Read, W., Undesirable and Fraudulent Behaviour in Online Auctions, In *Proceedings of the International Conference on Security and Cryptography (SECRYPT)*, pages 450–458, 2006.
- [10] Trevathan, J. and Read, W., A Simple Shill Bidding Agent, In *Proceedings of the 4th International Conference on Information Technology - New Generations (ITNG)*, to appear, 2007.
- [11] McCabe, A. and Trevathan, J., A New Approach to Avoiding the Local Extrema Trap, In *Proceedings of the 13th Computational Techniques and Applications Conference (CTAC)*, page 55 (1–10), 2006.
- [12] Trevathan, J. and McCabe, A., Pre-processing of On-line Signals in Noisy Environments, In *Proceedings of the 4th International Conference on Information Technology - New Generations (ITNG)*, pages 766–771, 2007.

- [13] Trevathan, J. and Read, W., Detecting Shill Bidding in Online English Auctions, *Social and Human Elements of Information Security: Emerging trends and counter measures*, Idea Group, in submission, 2006.
- [14] Trevathan, J. and Read, W., Detecting Collusive Shill Bidding, In *Proceedings of the 4th International Conference on Information Technology - New Generations (ITNG)*, pages 799–808, 2007.
- [15] Trevathan, J. and Read, W., A Software Architecture for Continuous Double Auctions, In *Proceedings of the International Association for Development of the Information Society (IADIS) – International Conference on Applied Computing*, pages 328–338, 2007.
- [16] Trevathan, J., An Adaptive Shill Bidding Agent, In *Proceedings of the International Conference on e-Business (ICE-B)*, to appear, 2007.
- [17] Trevathan, J., Privacy and Security Concerns for Online Share Trading, *IEEE Security and Privacy*, in submission, 2006.
- [18] Trevathan, J. and Read, W., Secure Online English Auctions, In *The 4th International Conference on E-Business and Telecommunications Engineers (ICETE) 2006 Best Papers Book*, Springer-Verlag, to appear, 2007.
- [19] Trevathan, J. and Read, W., Variable Quantity Market Clearing Algorithms, In *The 4th International Conference on E-Business and Telecommunications Engineers (ICETE) 2006 Best Papers Book*, Springer-Verlag, to appear, 2007.