Behavior Identification in Markets using Visualization and Network Analysis

Zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der Fakultät für
Wirtschaftswissenschaften
des Karlsruher Instituts für Technolgie (KIT)

genehmigte

DISSERTATION

von

Dipl.-Inform.Wirt. Hans Michael Blume

Tag der mündlichen Prüfung: 16.02.2012
Referent: Prof. Dr. Christof Weinhardt
Korreferent: Prof. Dr. Detlef Seese

Karlsruhe, 2011

Version of: February 17, 2012

Contents

List of Figures

List of Tables

Acknowledgement

I would like to take the opportunity to thank a number of people have generously given time, advice, encouragement and valuable information during the course of this research.

I would like to express my deep and sincere gratitude to my first supervisor, Professor Christof Weinhardt, Head of Chair of Information & Market Engineering, Institute of Information Systems and Management (IISM), Karlsruhe Institute of Technology (KIT). His wide knowledge and visionary thinking have been of great value for me. His understanding, encouraging and personal guidance have provided a good basis for the present thesis.

I am much in debt to my second supervisor, Professor Detlef Seese, Institute of Applied Informatics and Formal Description Methods, Karlsruhe Institute of Technology, for his detailed and constructive comments, and for his important support throughout this work.

Together they form part of the interdisciplinary doctoral program "Information Management and Market Engineering" at the Karlsruhe Institute of Technology funded by the DFG (Graduate School 895) which provided me a sound framework, funding, and background for my thesis. Everyone has helped me, in numerous discussions, to focus on the important questions.

Special thanks go to the STOCCER Team (funded by the BMBF in project 01HQ0522). Their fruitful cooperation supported this thesis in many ways. First of all, they provided me with a real data set making the research more interesting and challenging. Besides the data, we also worked closely together to exchange ideas and characteristics of fraud in sports prediction markets.

Prof. Stefan Seifert was a great help in discussing the scientific methods, proofreading, and in asking the right questions the right moment.

Furthermore, I want to thank the following persons for reviewing earlier versions of this manuscript (or parts of it) and for suggesting helpful improvements: Marc Adam, Arun Anandasivam, Jochen Stößer and Stefan Luckner, were so kind to improve in parts the English manuscript.

For the non-scientific side of my thesis, I particularly want to thank my parents for supporting me during all these years of university. Special thanks go to my wife Karmele Chavéz for being my eternal sunshine and being patient with me during all the weekends I spent writing. I also want to thank the IISM team and the IME colleagues with whom I spent a great time during the last years.

Theorem. *We cannot prove that a theory is true, but we can certainly show that a prediction is false. [following Karl Popper]*

1 Introduction

Online market platforms like eBay, Amazon, and Yahoo have a historically unique path to success. These companies started without branch offices, interacting with their customers exclusively online. Their success illustrates the importance electronic markets have achieved in our everyday life. The increasing penetration rate of high speed Internet access across the world combined with high costs for service staffing in first world countries promises high returns on online services with little or no branch offices involved.

Online services base their success on trust in infrastructure, mechanisms, and counter parties. The infrastructure is seen as reliable and trustworthy in many countries. But counter parties, such as service providers and —if involved— third parties are difficult to evaluate and more likely to defect. This holds especially true for electronic markets, where counter parties often use pseudonyms rather than their real names and may originate in other countries with different jurisdictions. The companies mentioned above try to foster trust via rating systems, reputation systems, and advertisement. However, recent numbers published by official institutions (National White Collar Crime Center 2006) about fraud and manipulation point towards a growing need for rule enforcement, monitoring, and prosecution tools for electronic market platforms.

People tend to manipulate markets in order to improve their stock to yield a higher benefit or to reduce uncertainty during trades. To assure a common understanding of the general terms such as *fraud*, *prediction markets,* and *market engineering* each term will be briefly introduced in the following subsections. Finally, an example will illustrate how manipulation can happen in an electronic sports prediction market and the difficulties encountered in discovering it. The example will pose the research questions addressed in the following chapters.

1.1 Market Design and Fraud

The term *fraud* is used in different contexts and with different meanings. The Concise Oxford Dictionary, for example, defines *fraud* as

> "criminal deception; the use of false representation to gain an unjust advantage."

Another definition is given by the Association of Certified Fraud Examiners :

> “the use of one’s occupation for personal enrichment through the deliberate misuse or application of the employing organization’s resources or assets.”
> (ACFE 2006)

In the example in the introduction, it is the violation of the general terms and conditions to improve the ranking of a certain account in order to increase the probability of winning the final lottery. This can be subsumed very well under the more general definition of the Concise Oxford Dictionary.

Following Darwin, it is natural to take advantage of others. It can be explained by our biological instincts. However, electronic markets theoretically provide a previously unreached level of transparency and rule enforcement via protocols that it should be possible to prevent manipulation by design. The questions become what fails on these markets and what can be improved in order to prevent or at least limit fraud? Electronic markets, such as STOCCER, Betfair, or the already mature eBay, are very popular. However, bad press and personal fraud experience affect Internet companies significantly since the cost for customers to switch to another platform are negligible.

> “Markets don’t always grow like weeds... some of them are hothouse orchids.”
> (Roth 2002)

This quotation illustrates that evolving markets do not always succeed. In fact, they might wilt very fast. Market Engineering addresses this problem by providing a framework to design, implement, test, and introduce markets. The underlying process has been formalized and introduced by Neumann (2004). The process steps are shown in figure 1.1.

The diagram shows the different steps a Market Engineer takes to design a new market. The term environmental analysis refers to the economic environment. In this step, questions like the following have to be answered: who are potential customers, what are their preferences, constraints, and endowments? Answering these questions creates a basic requirement list for designing the market.. The design phase includes both the conceptual design of the market as the design object and the embodiment design, wherein the abstract concepts develop into blueprints for protocols such as trading protocols. This is followed by the detail design and implementation phase. The detail design phase fills the gap between a blueprint and a concrete system design. After the first implementation, the incentive scheme is evaluated. This can be done via rough estimations, field studies, simulations, or experiments.

To facilitate the work of the market engineer, several tools have been proposed. While many tools set the focus on specific aspects such as simulations (e.g. *marketsim* (Porter et al. 2006)), experiments (e.g. *ztree* (Fischbacher 2007)) or even application fields (like *zocalo* on prediction markets (Hibbert 2005)), Meet2trade combines a complete set of tools supporting different steps of the market engineering process (Neumann et al. 2005; Weinhardt et al. 2006b). It has been successfully applied in various settings and contains

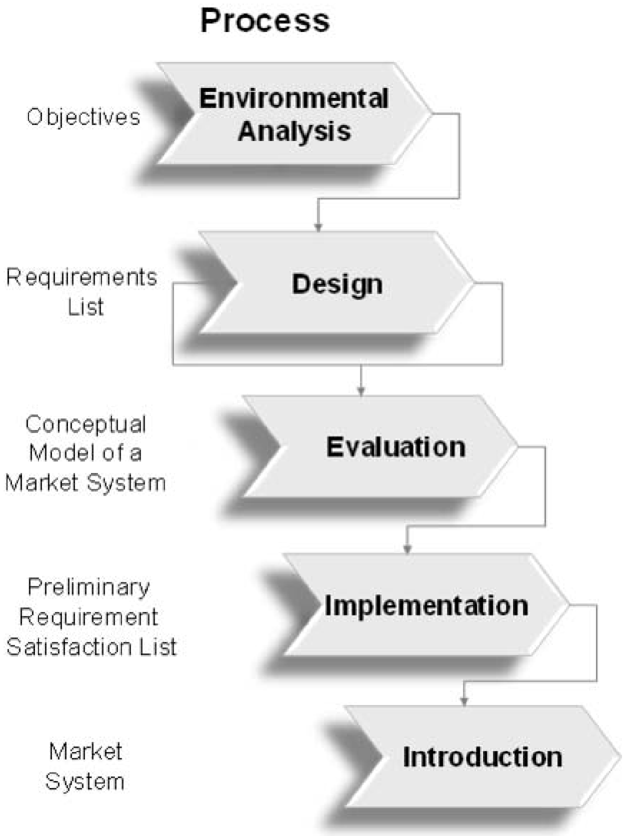

Figure 1.1: Market Engineering Process (Neumann 2004)

specific tools for market design (*MML* (Mäkiö and Weber 2005)), experiments (*MES* (Kolitz and Weinhardt 2006)) and agent based simulations (*AMASE* (Czernohous et al. 2003)). STOCCER was built on top of *meet2trade* (Weinhardt et al. 2005, 2006a) and the *psm* (Political Stock Market[1]).

In the age of the Internet, the time between an idea and its first publication is often minimized since being the first one offering a service can have a significant impact on your success (Andersen et al. 2003). This pressure to launch early naturally results in less time spent on the evaluation of how people may take advantage of the system. Time pressure and limited evaluation are two major factors influencing the shortcoming of the market engineering process in the Internet economy. Even when system abuse is evaluated, the electronic environment makes it hard to foresee every possible avenue of abuse. A bank robber will try to get a million dollars in a single attempt, but in the Internet it is easier to trick a million users to spend just a single dollar each (Cukier et al. 2007; August Bequai 2001). Things which are too much effort for a human being can be carried out by fully automated machines with little or no special programming knowledge at all.

Since errors in the incentive scheme and in system design can never be excluded, a market quality analysis has to be carried out on a regular basis after its introduction. This is the next phase in the market engineering process: *operation*. To facilitate the operation of a market, a network analysis based approach will be presented and applied to detect manipulation in sports prediction markets. The next subsection will introduce prediction markets in general and sports prediction markets in particular.

1.2 Prediction Markets

Prediction markets is only one of the terms used in scientific publications to describe this forecasting method. Other terms include information markets, decision markets, idea futures, forecasting markets, electronic markets and virtual stock markets. A recent literature overview comparing the different terms can be found in Tziralis and Tatsiopoulos (2007). Prediction markets are usually arranged very similar to common stock markets. This is also reflected in the names attributed to them. But unlike stock markets, whose primary concerns are market capitalization as well as resource and risk allocation, the primary goal of prediction markets is the aggregation of information(Luckner 2008).

Rhodes and Stumpf describe how market based forecasting methods have been used since the Middle Ages, and more recently, in presidential wagering markets in the United States (2004). With the broad availability of telecommunication and the Internet at the end of the 20th century, the application of prediction markets became easier than ever before. Nowadays, prediction markets are used in many fields including politics (Forsythe et al.

[1] Polticial Stock Market `http://psm.em.uni-karlsruhe.de`

1992), sports (Luckner et al. 2007), medicine (Polgreen et al. 2007), and entertainment (Pennock et al. 2000). Besides markets which are open to the public, prediction markets have also been employed within companies like Siemens or Hewlett-Packard in order to improve their decision making (Plott and Chen 2002; Ortner 1998).

The idea behind prediction markets is to translate the uncertain probability of future events into a tradeable contract. If an appropriate payoff scheme is applied, traders will reveal their true beliefs about the probability of the underlying event. According to the efficient market hypothesis (Fama 1970), the market will aggregate the diverse beliefs —reflected in the trading prices— into a single price (Hayek 1945). The aggregation continues as long as traders observe the current price and then trade against it until equilibrium is reached (McKelvey and Page 1990). The contracts have to accommodate the distinct types of predictions, for example, ranking of alternatives or yes-no questions. Wolfers and Zitzewitz distinguish three basic types of contracts (Wolfers 2004):

- *winner-take-all,* where a certain amount of money is paid only if the event occurs, otherwise nothing,
- *index,* the contract price is linked with a real number such as percentage of votes a candidate gains; and
- *spread,* a certain cutoff is defined such as whether a candidate gains more than a certain percentage of the votes.

In practice, variations of these contract forms are often employed. For example, if predicting a specific rank is desired, portfolio trading can be used in order to balance the different options. Portfolio trading is when a basket containing one of each contract is bought for the baseline price and then bought back from the market operator for the sum of the current market prices of all shares. This allows arbitrage traders to level the prices against the baseline assuming there is enough liquidity in the market.

The incentive scheme plays a key role in the motivation of the traders. Various schemes have been proposed and evaluated against each other. Some of the often used schemes are the publication of rankings performance based pay-offs and fixed pay-offs. In real-money markets the performance based pay-off is widely used. There is discussion regarding whether real-money markets are more accurate than play-money markets. However, recent findings indicate that they perform equally well (Servan-Schreiber et al. 2004). Because of the German gambling law (LottStV 2003), STOCCER used play-money. Every trader began with an initial endowment of 100,000 virtual currency units as well as 100 shares of each contract. In order to motivate traders, small prizes were raffled among the most active traders of the week as well as larger prizes among the final top 100 traders ranked by their portfolio.

The trading mechanism is the third key component of a prediction market — besides the contracts and the incentive scheme. Typical trading mechanisms are the continuous double auction (CDA), the call auction (CA), dynamic pari-mutuel markets (Pennock

2004) and market scoring rules (Hanson 2007). In STOCCER, both CDA and CA have been implemented, where the CDA had the higher trading activity as measured by the transactions per day (Geyer-Schulz et al. 2007).

The FIFA World Cup 2006 was held in Germany from June 9th to July 9th 2006. Thirty-two national teams qualified for the tournament. The tournament was organized in two stages: a group stage and a knock-out stage. All in all, 64 matches were played: 48 in the group stage and 16 in the knock-out stage. Overall, more than 1,700 traders registered on the platform and initiated on average more than 1,600 trades per day with a total of about 90,000 trades. The traders could trade on 19 different markets: 16 final rounds matches, 2 for the top goal scorers and a so-called championship market for shares of all 32 teams. The latter market was the one with the highest trading activity and will therefore be used as the data set for the following.

The championship market started on May 15th 2006 and ran until the end of the FIFA World Cup on July 9th 2006. The platform consisted of a trading system and a portal site to publish rankings, FAQs, news, and host discussion forums.

1.3 State of the Art in Fraud Detection

Fraud detection is investigated in different domains within the scientific community as well as industry. The solutions, metrics and results are of such a diverse nature, that several literature surveys have been published in the last years. Due to their extensive nature this section will discuss fraud detection in general and close with the critics and comments of the literature reviews. For further reading, the review by Bolton and Hand (2002) and of Phua et al. (2005) is especially recommended.

The literature regarding Fraud detection is highly related with statistics, data-mining, visualization and artificial intelligence. The methods are usually classified into supervised, semi-supervised and unsupervised. Application areas range from insurance fraud, mobile and telecommunication fraud, e-commerce fraud, on- and offline credit card fraud, identity fraud, and financial fraud – in markets as well as in accounting. Related areas are law enforcement, anti-terrorism systems, and intrusion detection systems. Several machine learning techniques need a pre-classified training data set. The term *labeled data* refers to such a data-set which usually is a real-world or generated data-set with an additional column containing the classification labels.

Focusing on neural networks and computational immune systems Margaret Weatherford (2002) published a short report about ongoing projects in fraud detection. According to her report, HNC Software applies back-propagation based neural networks with three layers of neurons. These are trained in supervised mode to detect credit card fraud in the Falcon Fraud Manager. A European research group is reported to use neural networks in mobile telecommunication fraud with an unsupervised learning approach, such that only non-fraudulent data is necessary. Finally, she presents a computational immune system used to detect fraud at Consignia, the former UK's Post Office Group.

Bolton and Hand (2002) published an overview covering several application areas and techniques. They distinguish between fraud prevention and fraud detection, the latter being continuously applied to find the cases where the first has failed. According to them, the interchange of ideas in the field of fraud detection is limited since fraudsters could use this information as well to circumvent the newly developed systems. Censored and unpublished data-sets restrict the evolution of the discipline as well. The large data sets of the typical application areas require the dection system to have a very high efficiency. Bolton and Hand see the uncertainty of the classifications as a fundamental point:

> "[...] we can seldom be certain, by statistical analysis alone, that a fraud has been perpetrated".

This means the only purposes of the fraud detection system is to distinguish a probability of the case via a suspicion score and then to alert the user. In the real world, fraud detection is a delicate topic since it is about trust. Neither news about fraud in certain companies, nor noting that you yourself are suspected to be fraudulent, will foster the necessary trust in the customer relationship.

After addressing fraud in general, Bolton and Hand discuss different tools for fraud detection and present selected publications from the following application areas: money laundering, network intrusion, credit card, telecommunications, medical, and scientific fraud. They conclude with a description of related areas such as insurance fraud, plagiarism and management fraud. Besides the previously mentioned essential importance of the detection speed, they consider the uncertainty about the false negatives to be key for the evaluation of the methods.

Kou et al. (2004) published a survey about fraud detection techniques focusing on credit card, telecommunications and computer intrusion. The survey is structured alongside these application areas and discusses within each of them the different techniques that have been applied so far and by whom. Besides the different publications about each technique, they describe open research issues they see in each of the application areas. Finally, they close with some criticisms. In credit card fraud detection, only few approaches are published in any detail. Among these, neural networks are very popular, but no data sets are available to implement or reproduce them. In the field of intrusion detection, data sets are available, but it is hard to reproduce or simulate realistic attack scenarios. Telecommunication fraud detection systems suffer from their inability to detect newly evolving fraud, which is not already present in the training data set and therefore requires system maintenance or an upgrade. It is especially important for the thresholds and parameters to be accurately defined.

An extensive review of the literature regarding the use of data mining techniques is presented by Phua et al. (2005). The mentioned criticism by Bolton and Hand regarding unavailable or unpublished data sets has a complete section dedicated to data and measurements. Phua et al. found 40 different sizes of data sets described in the literature, ranging from less than 500, up to 100 million, with a number of different attributes. The vast majority have less than 50 attributes with only 6 data sets containing more. None of the data sets are publicly available except for one relatively small automobile insurance data set.

Similar to the wide variety of data sets is the multitude of employed performance measures. The survey encountered as measures: explicit cost of fraud, misclassification cost, false positives, false negatives, real positives, entropy, conditional entropy, relative conditional entropy, information gain, detection time, Area under the Receiver Operating Curve (AUC), cross entropy (CXE), Brier score (mean squared error of prediction), Hellinger score, t-statistics, online vs. batch, number of different frauds detected and some problem specific criteria. Phua et al. regret that some recent studies still aim to only maximize accuracy (true positives vs. false positives) arguing that

> "in fraud detection, misclassification costs (false positive and false negative error costs) are unequal, uncertain, can differ from example to example, and can change over time".

They further state that

> "a false negative error is usually more costly than a false positive error".

In cases regarding markets this does not necessarily hold. Investigators from stock exchanges have stated in personal interviews with the author of this work that if a system comes up with several thousand cases/events/rule violations and most of them are false positives they stop believing in the system and therefore stop using it. Most of the investigating departments do not have the resources to manually verify all reported cases. Consequentially, if an automatic fraud detection tool shall support the investigators, it's *accuracy* is a core criterion. Depending on the application area, time to alarm may be of high importance as well. For example, online credit card fraud can cause high damage within minutes while insurance fraud detection may only require running a batch system overnight. In prediction markets, real-time detection is desirable but not indispensable.

Phua et al. (2005) not only review, but also criticize the current data-mining methods and techniques. According to them, research is often data oriented, while real-world applications are more resource and management dependent. Furthermore, the industry developed its own solutions independently, yet there has been no empirical evaluations of these commercial systems since the one by Abbott et al. (1998). Only 7 studies of the 51 claim to be implemented as actual fraud detection systems. Furthermore, very few are using temporal information and none are using spatial information. Researchers focus more on complex, non-linear supervised algorithms, but in the real world, due to missing labels and time restrictions, only semi-supervised and unsupervised approaches can be employed. Phua et al. (2005) predict that in the long run, faster and less complex algorithms will be winning. They close with the recommendation that data-mining should consider new approaches such as outlier detection (Hodge and Austin 2004), skewed classes (Weiss 2004), sampling (Domingos et al. 2002) or graph mining (Washio and Motoda 2003). These are applied in related application areas like law enforcement and intrusion detection.

After this brief overview of fraud detection in general, the next section will introduce fraud detection on prediction markets in particular.

1.4 Related Work on Fraud in Prediction Markets

While polls rely on the quality and independence of the underlying sample, markets are more complex and therefore the possibility of manipulation may raise concerns. The impact of manipulation on the performance of prediction markets is without a doubt an open question. Wolfers and Zitzewitz published a paper on five open questions about prediction markets; one of them is "How can markets limit manipulation?" (Wolfers and Zitzewitz 2006). Furthermore, they distinguish between the intent to manipulate the outcome which is to be predicted and the manipulation of the prediction market itself.

This distinction is also made by Ottaviani and Sørensen(2007). They refer to previous work in the financial literature of Allen and Gale (1992) and Vila (1989) for further classifications of manipulation. In their paper, they investigate agents' incentives to manipulate in corporate prediction markets using a theoretical model. They consider the posibility that agents may not be able to influence the price on the market but can influence the outcome –i.e. the project end or cost– by choosing a different effort level.

Some more research on manipulation has been conducted in the field of political stock markets (PSM). Hansen et al. (2004) discuss the effect of manipulation under the preconditions of indecisive voters and mass media coverage. They conclude that PSMs are vulnerable to manipulation and find that small contracts are especially effective at being manipulated. As a solution, they propose the reduction of market imperfections and the filtering of the prognosis (reduction of media coverage).

Bohm and Sonnegård (1999) tested a PSM in the context of a referendum, namely the Swedish referendum about joining the European Union. In this context, they evaluated the prediction quality of a PSM compared to polls and studied whether polls induce market activity. Furthermore, they introduced manipulation by adding a side competition. This competition collected a participation fee of 10 Swedish Krona from all participants payable to the account with the highest trading gains. This should not influence market accuracy as long as no coalitions are built. A coalition where n participants share the final benefit but ruin $n-1$ accounts in favor of one single account may distort prediction accuracy since their goal is now to transfer money from one account to another and not to strike the most likely market price. They conclude that unless the regular market is strong and has considerable economic power it is possible for any group to distort the prices at least for a certain period of time.

Most of the research on political stock markets consider outcome manipulation. In the field of market manipulation, Camerer (1998) reports some interesting observations. He observed that temporarily introduced manipulative bids in racetrack betting markets had a certain but statistically insignificant impact on market prices.

Hanson and Oprea (2004) take a contrary position. By modeling a market with a modification of the simplified Kyle (1985; 1989) model and adding a participant with a different price preference to the market, they find that the mean target price has no effect on prices and the price accuracy even increases. Their model has T participants of which one is a manipulator and N are informed traders (subset of T) who can obtain disclosed information about the bias of the manipulator and the real asset price v. In a follow up experiment, Oprea et al. (2006) investigated the influence of manipulation on price aggregation and on observers in a lab experiment. Therein manipulators did not succeed in raising the prices and failed in lowering them. Even though aggregation did not work properly, there was no statistically significant influence of manipulation on information aggregation. Dimitrov and Sami indicate in their recent publication (2008) that non-myopic strategies in prediction markets may actually be profitable and can temporarily distort prices.

This discussion of recent work on manipulation in prediction markets already indicates the early stage of the field. Hanson (2006) states that empirical results on price manipulation are mixed and the evidence is weak for actual manipulation in political stock markets. Moreover, a large part of the earlier work focuses on outcome manipulation where traders who care about the outcome try to influence prices. This is negligible in our field of sports prediction markets since participants of our market can most likely not influence the outcome of the predicted event such as FIFA World Cup 2006.

After a general introduction to the field, a detailed example is presented on how participants commited fraud on the sports prediction market platform STOCCER to illustrate the problems of revelation and raise the questions the following chapters will answer.

1.5 Fraud Examples in Sports Prediction Markets

On the sports prediction market platform STOCCER[2], people can trade their expectations in form of stocks about the results of various soccer events. The platform consists of a trading system, information portal, and a user forum.

By German law, it is forbidden to pay the participants according to the profits they achieve on the platform since the stock exchange is a zero sum game. This implies that people

[2] http://www.stoccer.com. The STOCCER project was funded by the German Federal Ministry for Education and Research under grant number 01HQ0522.

can loose as well which means it would be gambling for them. Gambling is a state-owned monopoly. In order to motivate people to trade on the platform during the FIFA World Cup 2006, prizes were raffled among the 100 most successful accounts. Participation was free.

A promising strategy in such an incentive scheme is to create an account for each team and play the make-or-break strategy. This means a participant sells all the shares of his account except those of team x and then uses the benefits from his selling to buy even further shares of team x, all under the belief that this will be the winning team. To reduce the risk of backing the wrong horse, the participant creates a separate account for each team. To prevent this behavior, the general terms and conditions excluded the possibility by permitting only one account per physical person. Violation was threatened with exclusion of the lottery. But in order to do so, it needed to be revealed in the first place.

During the FIFA World Cup 2006, some accounts moved into the top 100 ranking shortly after being created. Participants who had been on the platform for a while, started complaining in the online forums and questioning the market operator by email, asking how this was possible. Some even investigated on their own and located price drops or peaks in certain markets, publishing their suspicion in the online forum. The market operator started investigating —following the hints— and found evidence of collusion. For illustration purposes, one case is presented here in more detail. The screenshots are taken from the prototype, presented in Section 4. However, in 2006, the operator had only an SQL-query interface which made manual investigations very time consuming.

The transaction listing from the account *Soccer securities* in Figure 1.2 shows a three day period during the world cup. The pop-up menu in the lower part of the screenshot lists all matches in which Tunisia was playing. They lost the last match on June 23rd and dropped out of the competition. The final value of the Tunesia share was € 0 according to the pay-off rules of that market. Apparently *Soccer securities* was trying a make-or-break strategy on Tunisia. He was buying during the days before the last match all shares he could get by placing a large buy order at € 2.31. This presumption is supported by the structure of his portfolio (see Table 1.1). The table shows in the rows the different shares traded on the market and in the colums the number of shares (*Amount*) the account was holding. The column *Gain/Loss* is the profit of all sell transactions so far minus the costs of all buy transactions so far. The last column *Final Value* is the value of the shares when the market closed and the final pay-offs per team were known.

Even after Tunisia dropped out on the 23rd he continued buying shares indicated by the blue marked transactions. But then, between 11:30 am and 11:34 am, he changed his mind and sold all his Tunisia shares within less than half an hour to three accounts only for € 2.69, making more than € 280,000.00 profit and pushing himself from the last ranks right into the top 100 of the almost 2,000 participants.

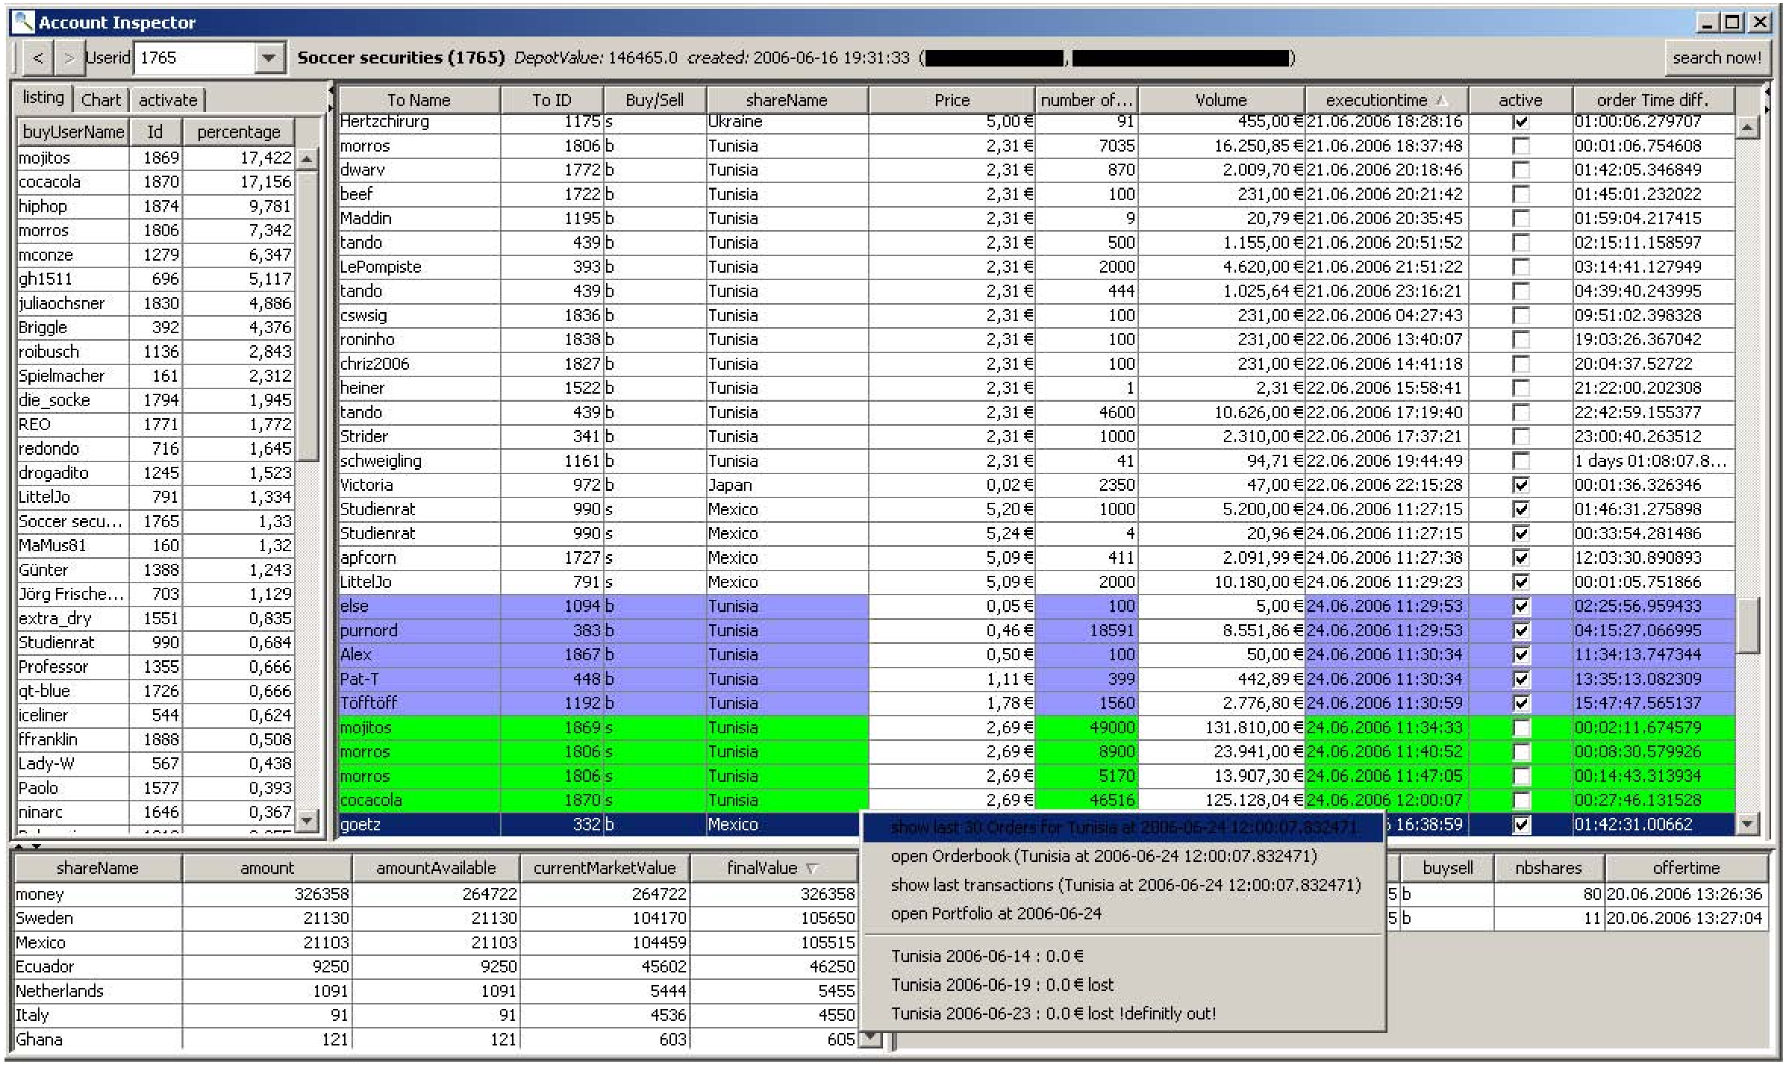

Account Inspector

Userid 1765 Soccer securities (1765) *DepotValue:* 146465.0 *created:* 2006-06-16 19:31:33 ([redacted], [redacted])

buyUserName	Id	percentage
mojitos	1869	17,422
cocacola	1870	17,156
hiphop	1874	9,781
morros	1806	7,342
mconze	1279	6,347
gh1511	696	5,117
juliaochsner	1830	4,886
Briggle	392	4,376
roibusch	1136	2,843
Spielmacher	161	2,312
die_socke	1794	1,945
REO	1771	1,772
redondo	716	1,645
drogadito	1245	1,523
LittelJo	791	1,334
Soccer secu...	1765	1,33
MaMus81	160	1,32
Günter	1388	1,243
Jörg Frische...	703	1,129
extra_dry	1551	0,835
Studienrat	990	0,684
Professor	1355	0,666
qt-blue	1726	0,666
iceliner	544	0,624
ffranklin	1888	0,508
Lady-W	567	0,438
Paolo	1577	0,393
ninarc	1646	0,367

To Name	To ID	Buy/Sell	shareName	Price	number of...	Volume	executiontime	active	order Time diff.
Hertzchirurg	1175	s	Ukraine	5,00 €	91	455,00 €	21.06.2006 18:28:16	☑	01:00:06.279707
morros	1806	b	Tunisia	2,31 €	7035	16.250,85 €	21.06.2006 18:37:48	☐	00:01:06.754608
dwarv	1772	b	Tunisia	2,31 €	870	2.009,70 €	21.06.2006 20:18:46	☐	01:42:05.346849
beef	1722	b	Tunisia	2,31 €	100	231,00 €	21.06.2006 20:21:42	☐	01:45:01.232022
Maddin	1195	b	Tunisia	2,31 €	9	20,79 €	21.06.2006 20:35:45	☐	01:59:04.217415
tando	439	b	Tunisia	2,31 €	500	1.155,00 €	21.06.2006 20:51:52	☐	02:15:11.158597
LePompiste	393	b	Tunisia	2,31 €	2000	4.620,00 €	21.06.2006 21:51:22	☐	03:14:41.127949
tando	439	b	Tunisia	2,31 €	444	1.025,64 €	21.06.2006 23:16:21	☐	04:39:40.243995
cswsig	1836	b	Tunisia	2,31 €	100	231,00 €	22.06.2006 04:27:43	☐	09:51:02.398328
roninho	1838	b	Tunisia	2,31 €	100	231,00 €	22.06.2006 13:40:07	☐	19:03:26.367042
chriz2006	1827	b	Tunisia	2,31 €	100	231,00 €	22.06.2006 14:41:18	☐	20:04:37.52722
heiner	1522	b	Tunisia	2,31 €	1	2,31 €	22.06.2006 15:58:41	☐	21:22:00.202308
tando	439	b	Tunisia	2,31 €	4600	10.626,00 €	22.06.2006 17:19:40	☐	22:42:59.155377
Strider	341	b	Tunisia	2,31 €	1000	2.310,00 €	22.06.2006 17:37:21	☐	23:00:40.263512
schweigling	1161	b	Tunisia	2,31 €	41	94,71 €	22.06.2006 19:44:49	☐	1 days 01:08:07.8...
Victoria	972	b	Japan	0,02 €	2350	47,00 €	22.06.2006 22:15:28	☑	00:01:36.326346
Studienrat	990	s	Mexico	5,20 €	1000	5.200,00 €	24.06.2006 11:27:15	☑	01:46:31.275898
Studienrat	990	s	Mexico	5,24 €	4	20,96 €	24.06.2006 11:27:15	☑	00:33:54.281486
apfcorn	1727	s	Mexico	5,09 €	411	2.091,99 €	24.06.2006 11:27:38	☑	12:03:30.890893
LittelJo	791	s	Mexico	5,09 €	2000	10.180,00 €	24.06.2006 11:29:23	☑	00:01:05.751866
else	1094	b	Tunisia	0,05 €	100	5,00 €	24.06.2006 11:29:53	☑	02:25:56.959433
purnord	383	b	Tunisia	0,46 €	18591	8.551,86 €	24.06.2006 11:29:53	☑	04:15:27.066995
Alex	1867	b	Tunisia	0,50 €	100	50,00 €	24.06.2006 11:30:34	☑	11:34:13.747344
Pat-T	448	b	Tunisia	1,11 €	399	442,89 €	24.06.2006 11:30:34	☑	13:35:13.082309
Töfftöff	1192	b	Tunisia	1,78 €	1560	2.776,80 €	24.06.2006 11:30:59	☑	15:47:47.565137
mojitos	1869	s	Tunisia	2,69 €	49000	131.810,00 €	24.06.2006 11:34:33	☐	00:02:11.674579
morros	1806	s	Tunisia	2,69 €	8900	23.941,00 €	24.06.2006 11:40:52	☐	00:08:30.579926
morros	1806	s	Tunisia	2,69 €	5170	13.907,30 €	24.06.2006 11:47:05	☐	00:14:43.313934
cocacola	1870	s	Tunisia	2,69 €	46516	125.128,04 €	24.06.2006 12:00:07	☐	00:27:46.131528
goetz	332	b	Mexico				5 16:38:59	☑	01:42:31.00662

show last 30 Orders for Tunisia at 2006-06-24 12:00:07.832471
open Orderbook (Tunisia at 2006-06-24 12:00:07.832471)
show last transactions (Tunisia at 2006-06-24 12:00:07.832471)
open Portfolio at 2006-06-24

Tunisia 2006-06-14 : 0.0 €
Tunisia 2006-06-19 : 0.0 € lost
Tunisia 2006-06-23 : 0.0 € lost !definitly out!

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	326358	264722	264722	326358
Sweden	21130	21130	104170	105650
Mexico	21103	21103	104459	105515
Ecuador	9250	9250	45602	46250
Netherlands	1091	1091	5444	5455
Italy	91	91	4536	4550
Ghana	121	121	603	605

	buysell	nbshares	offertime
5	b	80	20.06.2006 13:26:36
5	b	11	20.06.2006 13:27:04

The system automatically highlights disadvantagous transactions in blue, advantagous in green.

Figure 1.2: Transactions of the account *Soccer securities*, 21-24 of June, 2006

Table 1.1: Portfolio balance for *Soccer securities* before and after the fortunate transactions

	Portfolio at 11:31:00		**Portfolio 12:01:00**		
Share	**Amount**	**Gain/Loss**	**Amount**	**Gain/Loss**	**Final Value**
Italy	91	1150	91	1150	4550
Mexico	483	-2512	483	-2512	2415
Netherlands	91	1145	91	1145	455
Tunisia	***109586***	***-157006***	***0***	***137779***	***0***
Japan	22541	-12047	22541	-12047	0
Paraguay	3591	-358	3591	-358	0
Iran	3165	-1372	3165	-1372	0
Korea Republic	191	0	191	0	0
Saudi Arabia	191	0	191	0	0
Poland	91	200	91	200	0
Côte d Ivoire	91	1	91	1	0
Czech Republic	45	1142	45	1142	0
Serbia and Montenegro	0	1174	0	1174	0
Spain	0	2346	0	2346	0
Sweden	0	1126	0	1126	0
Switzerland	0	1457	0	1457	0
Togo	0	3857	0	3857	0
Trinidad and Tobago	0	166	0	166	0
Ukraine	0	971	0	971	0
Angola	0	4557	0	4557	0
USA	0	210	0	210	0
Argentina	0	4307	0	4307	0
Australia	0	887	0	887	0
Brazil	0	5711	0	5711	0
Costa Rica	0	105	0	105	0
Croatia	0	870	0	870	0
Ecuador	0	1222	0	1222	0
England	0	2514	0	2514	0
France	0	-14684	0	-14684	0
Germany	0	3858	0	3858	0
Ghana	0	60753	0	60753	0
Portugal	0	2109	0	2109	0
SUM		**-86141**		**208644**	

If the change of luck is not suspicious enough, a glance at the counter parties provides more insight. Figure 1.3 on the facing page shows the three accounts who bought the worthless Tunisia shares from *Soccer securities*. Among these three accounts, the account *morros* is particularly interesting. He has a longer trading history and apparently also played a make-or-break strategy on two outsiders: Tunisia and Mexico. Note, that after a sequence of Mexico acquisitions between € 5.40 and € 5.67, he is selling them for € 5.36 to *Soccer securities*. Both the acquisition the disposition took place on June 21 right after the end of the game where Mexico won against Portugal and thereby qualified for the next round. The shares of all teams of the second round paid a minimum of € 5.0. For Tunisia, *morros* exhibits less ambition but still buys large amounts here and there from *Soccer securities*. Therefore, *Soccer securities* can be seen as the major buy and sell partner of *morros* in terms of turnover.

Even more obvious are the cases of the accounts *mojitos* and *cocacola*, where *Soccer securities* is the only trading partner and the two accounts were created on the same day just two hours earlier. The whole pattern of both accounts is very similar and it is conjecturable that they are the same persons. Since the only account that benefited from their actions was *Soccer securities*, the market operators excluded *Soccer securities*, *cocacola* and *mojitos*, but failed to notice the relation to *morros*.

The operator considered these fraudsters harmful since their trading behavior, at least temporarily, influenced the market prices and thus also the prediction accuracy of the market. Furthermore, complaints which were published in the discussion forum about fraudsters in the market, such as "Is it possible to exclude cheaters from the market? They screw up the market" or "My first experience in STOCCER is that some traders are cheating", exemplify the traders' annoyance. In the worst case, well-informed and rather motivated traders may stop participating in the market. Consequently, it is desirable to prevent fraud in prediction markets.

Manipulation in prediction markets is not only limited to STOCCER. Betfair, one of the largest commercial online betting platforms for sport events had to void a market for a Polish 4th tier tennis match involving No. 4-ranked Nikolay Davydenko of Russia and No. 87-ranked Martin Vassallo Arguello of Argentina (Culpepper 2007). Such a match usually does not attract too much attention. But even before the match began, the odds went away from the logical favorite toward the underdog. The total was £ 3.59 million, ten times higher than usual. Though the favorite won the first set, he lost the second set and finally dropped off in the third set with a foot injury. Betfair and ATP started investigations. Betfair started offering to share their data with sports governing bodies in order to reveal corruption. But many of them —including the International Olympic Committee— refused the offer (Drape 2008). The Betfair case illustrates what can happen if a group of people collude to increase their profit. When the sportsmen are involved, it is even worse since this damages the sport's reputation, like recent cases in European Soccer Matches illustrate.

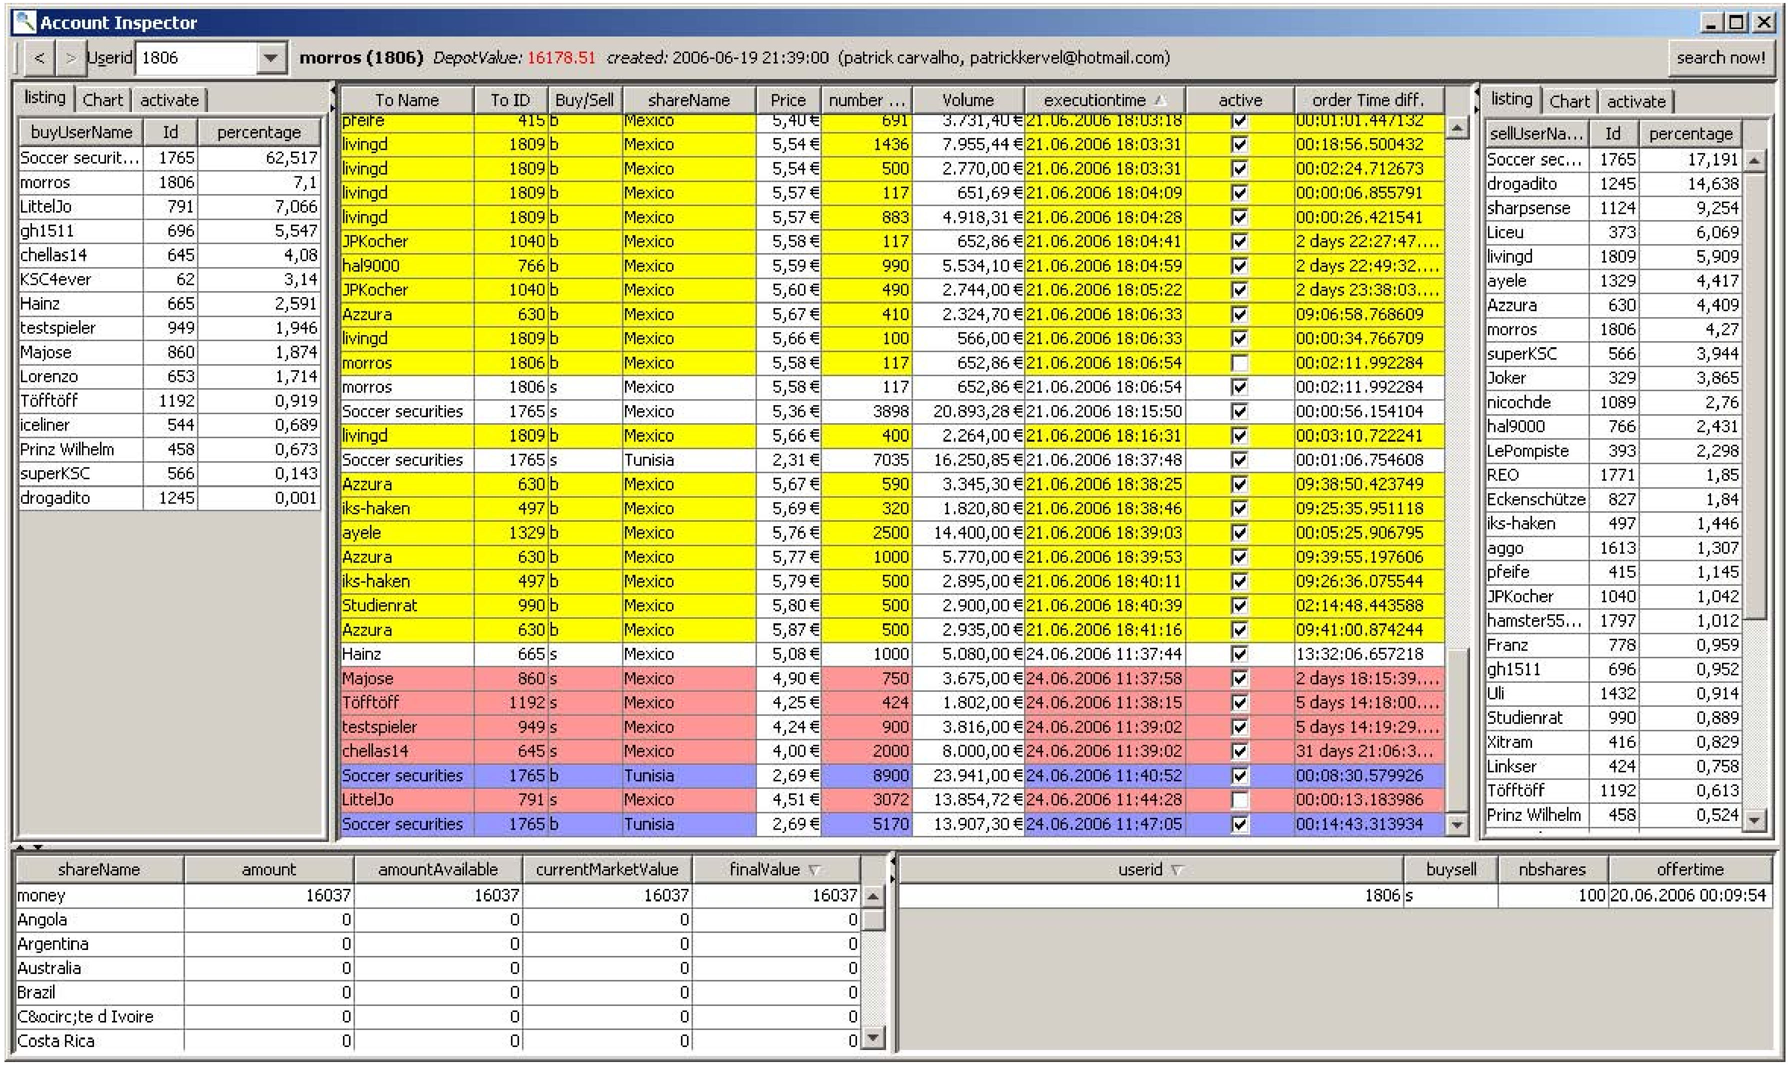

morros (1806) *DepotValue:* 16178.51 *created:* 2006-06-19 21:39:00 (patrick carvalho, patrickkervel@hotmail.com)

buyUserName	Id	percentage
Soccer securit…	1765	62,517
morros	1806	7,1
LittelJo	791	7,066
gh1511	696	5,547
chellas14	645	4,08
KSC4ever	62	3,14
Hainz	665	2,591
testspieler	949	1,946
Majose	860	1,874
Lorenzo	653	1,714
Töfftöff	1192	0,919
iceliner	544	0,689
Prinz Wilhelm	458	0,673
superKSC	566	0,143
drogadito	1245	0,001

To Name	To ID	Buy/Sell	shareName	Price	number …	Volume	executiontime	active	order Time diff.
pfeife	415	b	Mexico	5,40 €	691	3.731,40 €	21.06.2006 18:03:18	☑	00:01:01.447132
livingd	1809	b	Mexico	5,54 €	1436	7.955,44 €	21.06.2006 18:03:31	☑	00:18:56.500432
livingd	1809	b	Mexico	5,54 €	500	2.770,00 €	21.06.2006 18:03:31	☑	00:02:24.712673
livingd	1809	b	Mexico	5,57 €	117	651,69 €	21.06.2006 18:04:09	☑	00:00:06.855791
livingd	1809	b	Mexico	5,57 €	883	4.918,31 €	21.06.2006 18:04:28	☑	00:00:26.421541
JPKocher	1040	b	Mexico	5,58 €	117	652,86 €	21.06.2006 18:04:41	☑	2 days 22:27:47.…
hal9000	766	b	Mexico	5,59 €	990	5.534,10 €	21.06.2006 18:04:59	☑	2 days 22:49:32.…
JPKocher	1040	b	Mexico	5,60 €	490	2.744,00 €	21.06.2006 18:05:22	☑	2 days 23:38:03.…
Azzura	630	b	Mexico	5,67 €	410	2.324,70 €	21.06.2006 18:06:33	☑	09:06:58.768609
livingd	1809	b	Mexico	5,66 €	100	566,00 €	21.06.2006 18:06:33	☑	00:00:34.766709
morros	1806	b	Mexico	5,58 €	117	652,86 €	21.06.2006 18:06:54	☐	00:02:11.992284
morros	1806	s	Mexico	5,58 €	117	652,86 €	21.06.2006 18:06:54	☑	00:02:11.992284
Soccer securities	1765	s	Mexico	5,36 €	3898	20.893,28 €	21.06.2006 18:15:50	☑	00:00:56.154104
livingd	1809	b	Mexico	5,66 €	400	2.264,00 €	21.06.2006 18:16:31	☑	00:03:10.722241
Soccer securities	1765	s	Tunisia	2,31 €	7035	16.250,85 €	21.06.2006 18:37:48	☑	00:01:06.754608
Azzura	630	b	Mexico	5,67 €	590	3.345,30 €	21.06.2006 18:38:25	☑	09:38:50.423749
iks-haken	497	b	Mexico	5,69 €	320	1.820,80 €	21.06.2006 18:38:46	☑	09:25:35.951118
ayele	1329	b	Mexico	5,76 €	2500	14.400,00 €	21.06.2006 18:39:03	☑	00:05:25.906795
Azzura	630	b	Mexico	5,77 €	1000	5.770,00 €	21.06.2006 18:39:53	☑	09:39:55.197606
iks-haken	497	b	Mexico	5,79 €	500	2.895,00 €	21.06.2006 18:40:11	☑	09:26:36.075544
Studienrat	990	b	Mexico	5,80 €	500	2.900,00 €	21.06.2006 18:40:39	☑	02:14:48.443588
Azzura	630	b	Mexico	5,87 €	500	2.935,00 €	21.06.2006 18:41:16	☑	09:41:00.874244
Hainz	665	s	Mexico	5,08 €	1000	5.080,00 €	24.06.2006 11:37:44	☑	13:32:06.657218
Majose	860	s	Mexico	4,90 €	750	3.675,00 €	24.06.2006 11:37:58	☑	2 days 18:15:39.…
Töfftöff	1192	s	Mexico	4,25 €	424	1.802,00 €	24.06.2006 11:38:15	☑	5 days 14:18:00.…
testspieler	949	s	Mexico	4,24 €	900	3.816,00 €	24.06.2006 11:39:02	☑	5 days 14:19:29.…
chellas14	645	s	Mexico	4,00 €	2000	8.000,00 €	24.06.2006 11:39:02	☑	31 days 21:06:3…
Soccer securities	1765	b	Tunisia	2,69 €	8900	23.941,00 €	24.06.2006 11:40:52	☑	00:08:30.579926
LittelJo	791	s	Mexico	4,51 €	3072	13.854,72 €	24.06.2006 11:44:28	☐	00:00:13.183986
Soccer securities	1765	b	Tunisia	2,69 €	5170	13.907,30 €	24.06.2006 11:47:05	☑	00:14:43.313934

sellUserNa…	Id	percentage
Soccer sec…	1765	17,191
drogadito	1245	14,638
sharpsense	1124	9,254
Liceu	373	6,069
livingd	1809	5,909
ayele	1329	4,417
Azzura	630	4,409
morros	1806	4,27
superKSC	566	3,944
Joker	329	3,865
nicochde	1089	2,76
hal9000	766	2,431
LePompiste	393	2,298
REO	1771	1,85
Eckenschütze	827	1,84
iks-haken	497	1,446
aggo	1613	1,307
pfeife	415	1,145
JPKocher	1040	1,042
hamster55…	1797	1,012
Franz	778	0,959
gh1511	696	0,952
Uli	1432	0,914
Studienrat	990	0,889
Xitram	416	0,829
Linkser	424	0,758
Töfftöff	1192	0,613
Prinz Wilhelm	458	0,524

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	16037	16037	16037	16037
Angola	0	0	0	0
Argentina	0	0	0	0
Australia	0	0	0	0
Brazil	0	0	0	0
C&ocirc;te d Ivoire	0	0	0	0
Costa Rica	0	0	0	0

userid	buysell	nbshares	offertime
1806	s	100	20.06.2006 00:09:54

(a) *morros*

Figure 1.3: Transactions of the collusion partners of *Soccer securities*

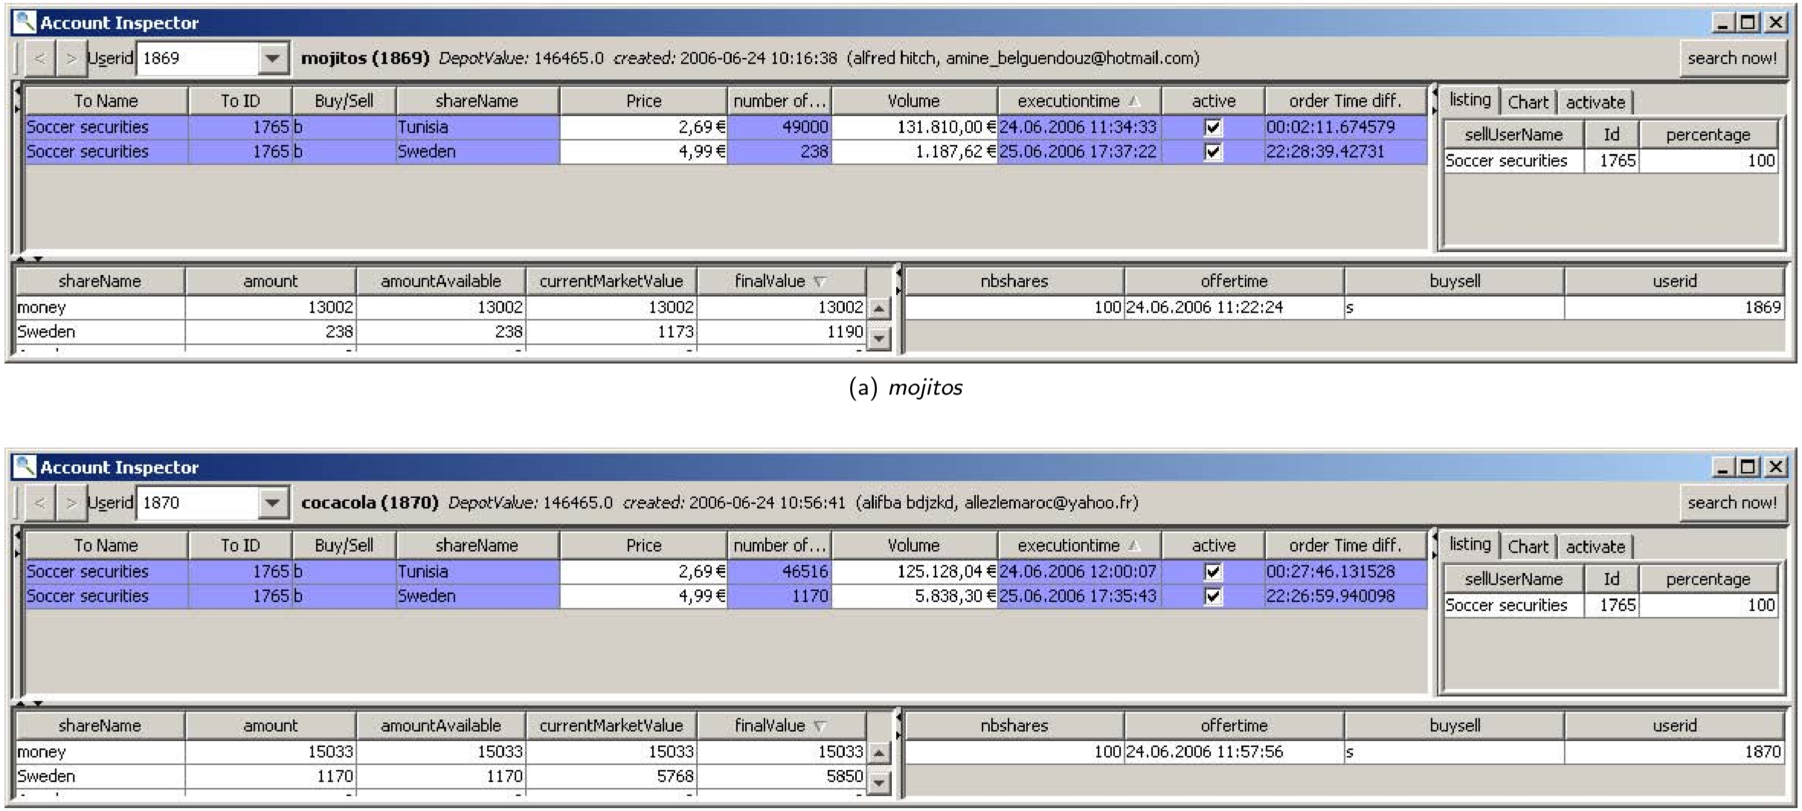

1869 **mojitos (1869)** *DepotValue:* 146465.0 *created:* 2006-06-24 10:16:38 (alfred hitch, amine_belguendouz@hotmail.com)

To Name	To ID	Buy/Sell	shareName	Price	number of...	Volume	executiontime	active	order Time diff.
Soccer securities	1765	b	Tunisia	2,69 €	49000	131.810,00 €	24.06.2006 11:34:33	☑	00:02:11.674579
Soccer securities	1765	b	Sweden	4,99 €	238	1.187,62 €	25.06.2006 17:37:22	☑	22:28:39.42731

sellUserName	Id	percentage
Soccer securities	1765	100

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	13002	13002	13002	13002
Sweden	238	238	1173	1190

nbshares	offertime	buysell	userid
100	24.06.2006 11:22:24	s	1869

(a) *mojitos*

1870 **cocacola (1870)** *DepotValue:* 146465.0 *created:* 2006-06-24 10:56:41 (alifba bdjzkd, allezlemaroc@yahoo.fr)

To Name	To ID	Buy/Sell	shareName	Price	number of...	Volume	executiontime	active	order Time diff.
Soccer securities	1765	b	Tunisia	2,69 €	46516	125.128,04 €	24.06.2006 12:00:07	☑	00:27:46.131528
Soccer securities	1765	b	Sweden	4,99 €	1170	5.838,30 €	25.06.2006 17:35:43	☑	22:26:59.940098

sellUserName	Id	percentage
Soccer securities	1765	100

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	15033	15033	15033	15033
Sweden	1170	1170	5768	5850

nbshares	offertime	buysell	userid
100	24.06.2006 11:57:56	s	1870

(b) *cocacola*

Figure 1.4: Transactions of the collusion partners of *Soccer securities* continued

Manipulations exist on political stock markets as well. Rhode and Strumpf documented in their publications (2004; 2007) several cases on different markets throughout the last 80 years. The manipulations they address are more concerned with price manipulation and its consequences. However, none of the cases succeeded in the long run. On historical Wall Street Betting Markets manipulators have changed the prices during several days, but they have only succeeded for several minutes on recent Tradesports[3] markets.

In general, fraud in electronic markets increased during the last few years and has become a serious problem for electronic market platforms such as eBay or Amazon. The Internet Crime Complaint Center received 207,492 complaint submissions in 2006 (National White Collar Crime Center 2006). Of those, 86,279 complaints were referred for further investigation with a total dollar loss of $198.44 million and a median dollar loss of $724.00 per complaint. Almost half of them (44.9%) have been Internet auction fraud. In 2003, a report from the SIRCA (2003) states a total loss due to identity fraud of $1 billion dollars just in Australia. Recent estimations in money laundering see it as the third biggest business in the world (Robinson 1998). These figures illustrate the economic dimensions and the growing need to address the underlying problems of market manipulation and fraud.

All these different types of manipulation and fraud have one thing in common: a group of persons or a one person with multiple identities join forces to reap benefits either inside the market or in a super game (e.g. in case of STOCCER the final lottery, in political stock markets the elections, etc.). Several identities collude to manipulate the market or game. Thus the main questions addressed in the following chapters are:

1. **Can the social structure between market participants be revealed by analysing only transaction data?**
2. **How can the social structure be used to detect collusion?**

The first question is clearly limited to markets where participants are trading with each other without a central intermediary such as a market or book maker. If an intermediary is present, the group of manipulators has to either include the intermediary or work around the other individuals decisions.

1.6 Summary

The recent rise of fraud and manipulation in electronic markets in general and in prediction markets in particular is a key problem for trust in current market platforms. Two real world examples have been presented to illustrate different aspects of manipulation. The majority of the scientific approaches to detect fraud use complex, supervised methods like neural networks and depends on labeled data sets. Disagreement on a common metric,

[3] Tradesports `http://www.tradesports.com/`

unpublished data sets, and the problem of the closely following fraudsters hinders further development of the field.

Since the trustworthiness of an electronic market, be it a commercial or a prediction market, is the basis for the customers to use the platform; bad reputation and manipulation is a serious threat to the operators. Most of the operators do not have the staff or the capacity for complete transaction monitoring. So the questions to be answered in the following chapters are:

1. Can the social structure between market participants be revealed by analysing only transaction data?
2. How can the social structure be used to detect collusion?

The chapters are organized as follows: The second chapter will present two different detection heuristics. The third chapter will propose a visualization which can be used to monitor a market. Chapter four describes a prototype implementation which is followed by an evaluation in chapter five. Chapter six closes with a summary and outlook.

2 Trading Pattern Detection

Trading patterns play a key role in fraud detection. Market designers can use them as a feedback to their design parameters, and traders could use them to develop counter-strategies.

Pattern recognition in fraud detection is a straight forward approach, in order to generate warnings to authorities if previously investigated crime patterns appear again. Though the term *pattern recognition* has a rather narrow meaning in computer science, there are a variety of data-mining algorithms for this purpose. For example, event-based systems raise events on certain predefined conditions or patterns and rule-based systems have hard coded rules for each pattern. Only a few approaches try new, uncommon techniques. Some examples are the work of Hill, who looks for statistically uncommon occurrences of first digits (Hill 1995) in accounting, game theoretical modelling of fraud strategies and development of counter strategies (Dalvi et al. 2004) and logical rule improvement by Kim et al. (2003). The problem of patterns is that they are very short-lived. Once the fraudsters become aware of a search pattern which matches their strategy, they alter their patterns in order to stay below the thresholds or deviate from the pattern and thus create a new one. Therefore operators have to continuously fine-tune the thresholds and update the pattern database. Consequently, a flexible pattern that automatically adapts to the current situation without any fixed thresholds and a system the automatically learns new patterns promises the highest success against fraudsters.

This flexibility is provided by artificial intelligence (AI) algorithms. AI research has applied adaptive algorithms like neural networks, genetic algorithms and case based reasoning to fraud detection. The problem of many of these AI approaches is that they act as a black box. The reason why a case had been classified as fraud because of a neural network is hard to comprehend. Without the reason for the classification, the investigator has to analyze the case in depth without knowing what to look for. Comments in a pattern database can at least link to previously judged cases and a rule based system can mark the element, which violates the rule. Thus, core characteristics of a good fraud detection system are *comprehensibility* and *flexibility*. Another problem of machine learning algorithms is that they need an accurately classified training data set. In several fields of fraud detection like money-laundering or telecommunications, the exact classification is often impossible. Accounts may turn out to be fraudsters several month or years later, as shown in Bolton and Hand (2002). Furthermore, trained algorithms can only recognize existing fraud patterns and are less flexible for new evolving patterns.

A very important question in the design of a fraud detection system is which performance measure to apply. For example, one could minimize false positives or false negatives. The SAS Institute stated in 1996, that the objective of fraud detection is to maximize correct predictions and to maintain incorrect predictions at an acceptable level (SAS Institute, 1996). Therefore, false positives and false negatives have to be minimized at the same time. Bolton and Hand (2002) provide a numeric example, showing that even with an accuracy rate of 99% for the classification of a fraud case, 99% accuracy for a legitimate case, and an actual fraud rate of 0.1 percent, on average only 9 of out of 100 flagged cases will actually be fraud[1]. Considering the costs for an in depth analysis for all hundred cases, the system would still generate a significant overhead. Many other measures have been suggested for the different algorithms and domains (c.f. previous chapter). In the insurance industry, the comparison against manual evaluation is done quite often (von Altrock 1997; Brockett et al. 2002; Stefano and Gisella 2001; Belhadji et al. 2000). As the discussion in the previous chapter about fraud in prediction markets indicates, a certain amount of noise trading is even desirable for a market. Thus a hundred percent clean market is neither achievable nor desirable. But a very high degree of noise may harm the information aggregation and thereby the quality of the market and threaten legitimate users.

The best solution would detect the fraudsters by their intentions rather than by their transactions. Intentions, the very inside of our thoughts, cannot be revealed with the information stored in today's transaction systems. Though neuro-science is advancing lately, the intentions behind a transaction is still left to the interpretation of the human investigator. Phua et al. (2005) point out that new research topics in the field of anti-terrorism and law enforcement such as link analysis and graph mining should be taken into consideration for fraud detection as well. Bolton and Hand recommend Peer Group analysis to monitor inter-account behavior over time and suggest Break Point Analysis to monitor intra-account behavior over time. Also, using further information like the social context of a trader and the context of the transaction may help to improve classification quality. A social network analysis related approach has recently been presented in law enforcement (Wang et al. 2006).

In this chapter, two graph based pattern detection algorithms are going to be presented. The patterns originate from fraud detection but may be of interest for other applications as well. Analyzing social graphs using specific metrics like betweenness, centrality and others has already been suggested earlier. A good overview can be found in Wasserman and Faust (1994).

[1] 100 cases get flagged, if the algorithm evaluates a 10,010 records while recognizing 99% of the legimate users as legimate. With a fraud rate of 0.1 percent, 10,010 records would contain 10 real fraud cases. Though the fraud algorithm will recognize most of the fraud records properly, it still leaves the user 90-91 records to revise.

2.1 Related Work

After the very broad introduction into fraud detection in the previous chapter only publications closely related to the algorithms of this thesis will be addressed in this chapter.

Cortes et al. (2001) introduced graph based fraud detection in telecommunication; where fraudsters make calls for free using other customer's accounts. They extended the existing ideas of analyzing graph patterns to identify interesting nodes (e.g. Wasserman and Faust 1994; Kleinberg 1998) from static graphs to dynamic graphs, introducing a new data structure called communities of interest (COI). This dynamic graph data structure, which they later refined (Cortes et al. 2003), allowed them to slightly improve the manual classification in a telecommunications data set. The idea is that fraudsters tend to learn from other fraudsters. This implies a slightly higher probability for other fraudsters in their social community. They observed that the shortest path from a fraudster to any other fraudster is usually shorter than from a normal node to any fraudster. The disadvantage of this approach is that it needs some previously classified fraud cases to indicate further accounts for investigation. An investigator – knowing this detail – may even automatically revise the social neighborhood of each classified fraudster. Nevertheless, the work indicates that a dynamic transaction graph contains structures which may reveal fraudsters by their patterns.

The next logical step would be the analysis of patterns in transaction graphs to distinguish between normal and fraudulent nodes. This leads to the evolving area of graph mining. From the field of bio-informatics and chemistry, several graph mining algorithms are available (Inokuchi et al. 2003). However, general sub-graph isomorphism is a NP-complete problem (Garey and Johnson 1979) and Washio and Motada (2003) point out in a recent survey that there are still many open problems in this young but promising data mining field. Due to the early stage of the field, only a few graph mining implementations are available. Since molecular structures are large in number but rather small in size, the run-time of the mining algorithm is still acceptable for bio/chemistry applications. But market graphs grow easily beyond common molecular dimensions such as in the case of traders of Apple Inc. shares on public stock exchanges. Therefore, unspecific, frequent, sub-graph mining is still too slow for markets.

A less pattern-, but more graph-oriented, approach in fraud detection is link analysis. Lee et al. (1999) define *link analysis* as follows:

> [it] "determines relations between fields in the database records. Correlations of system features in audit data, for example, the correlation between command and argument in the shell command history data of a user, can serve as the basis for constructing normal usage profiles."

Goldberg et al. (1995; 1998) have already proposed using link analysis to build profiles for the revelation of money laundering. The links were made searchable through a query interface and displayed for further manual investigation using the NetMap tool (Davidson 1993). Though they argue that the power of the system is the "man in the loop"[2], they also have to admit that manual data analysis is not manageable any more.

Closely related to link analysis is *record linkage,* also known as the merge-purge-problem (Hernandez and Stolfo 1995), object isomerism (Chen et al. 1996) or instance identification (Wang and Madnick 1989). It refers to the mapping of identities in different data sources. An example could be that "Eick, 2009", "S. Eick" and "Stephen G. Eick" all refer to the same person in a citation data set. Mapping this ambiguous reference correctly is a common problem in information integration and of manifold nature. Records may have different keys, missing data, inaccurate information such as only estimates of the age of a criminal and data entry errors like misspellings or typing mistakes. The goal is to find duplicates within a database or corresponding records between different data sources. Identical duplicates can be found easily by sorting the database and looking for a sequence of identical records. But when relaxing the assumption from identical to only similar records, each record has to be evaluated against every other record in the worst case using $\frac{n*(n-1)}{2}$ comparisons. The solution to the explosion of comparisons is to first search and select suitable records for comparison followed by matching pairs with a distance measure. Several heuristics have been proposed to tackle the problem (e.g. Monge 2000; Elfeky et al. 2002). But none of them is suitable, if a person deliberately hides his identity or even uses a fictitious one.

This type of fraud is called *identity fraud* and has recently gained more attention (Office 2002; SIRCA 2003). The Home Office Identity Fraud Steering Committee, a group of governmental and associated organizations in the UK, defines identity fraud as follows:

> "Identity fraud occurs when a false identity or someone else's identity details are used to support unlawful activity, or when someone avoids obligation/liability by falsely claiming that he/she was the victim of identity fraud".

They define a *false identity* as:

> "(a) a fictitious (i.e. invented) identity; or (b) an existing (i.e. genuine) identity that has been altered to create a fictitious identity".

Since identity fraud often leads to criminal investigations, law enforcement is especially looking for algorithmic support. An interesting approach has recently been presented by Wang et al. (2006). They use social contextual information to match criminal identities. Since social contextual information is usually not stored in law enforcement databases, these features have to be extracted from personal records and crime incident reports. The obtained information is used to span the social network. They tested the four features: structural similarity, relational-, group- and personal distance, to train a C4.5

[2] This expression refers to the integration of human beings in the process of recognizing fraud. The information is gathered by the software and presented to the human investigator who analyzes and finally classifies the presented case.

decision tree algorithm and achieved a significant improvement of the prediction quality about which accounts belong to the same identity. C4.5 is a standard machine learning algorithm introduced 20 years earlier by Quinlan (1993). A related approach, again using a C4.5 decision tree algorithm to classify fraudsters on eBay, has been proposed by Chau et al. (2006). The classifier is based on features extracted from the account's transaction history and just provides a preliminary classification. In a second step, the social network structure is built from the bidding and transaction relationships and used to refine the classifications with a hidden markov model using a belief propagation algorithm.

On prediction markets Schröder used graph based analysis using eigensystem analysis techniques to detect particular trading behaviors (2009). He identified specific eigenvector compositions for particular trading patterns related to the ones described in this thesis. However, no information about precision and recall of this method is available. Also the eigensystem analysis is very computing intensive already for mid-size graphes ($O(|V|^2|E|)$ without any optimization).

Maranzato et al. have made tests to evaluate and verify the ratings given by social networks embedded in Latin American eBay competitor TodaOferta.com (2010). They used various kinds of activities (account creation, bidding, offering goods, accessing listings, etc.) to evaluate trustworthiness of ratings. After identifiying 17 indicators where activities happened from the same machine or IP address, they run a logistic regression to rank and weight the different indicators. The goal is to develop a tool to verify members ratings to foster trust into the right accounts and the marktet platform. In a previous publication, Maranzato et al. build the graph of who rated whom in oder to look for transitive rating relations and by this networks of similar rating accounts (2009).

2.2 Design of Detection Algorithms

Inspired by the work of Cortes et al. (2001) with respect to the shorter path lengths between two or more fraudsters than between legitimate accounts and fraudsters; a similar test was applied to the FIFA prediction market data set. Cortes et al. observed in telecommunication networks fraudsters finding ways to abuse the system to make calls on behalf of other or non existing customers. Being happy about their success they often shared their techniques with family or friends. Thus in the call graph a community of fraudsters appeared usually together. This behavior manifests itself in shorter paths between fraudsters than from normal users to fraudsters.

In STOCCER the manually discovered fraudsters were used for the labelling. While Cortes et al. saw the reason for the phenomenon in information dissemination of the current fraud techniques, in markets it is more the endowment of a ruined account flowing to the winning main account of the fraudster. Since each share can be seen as a separate market and the observed fraud techniques did not span different shares, a graph for

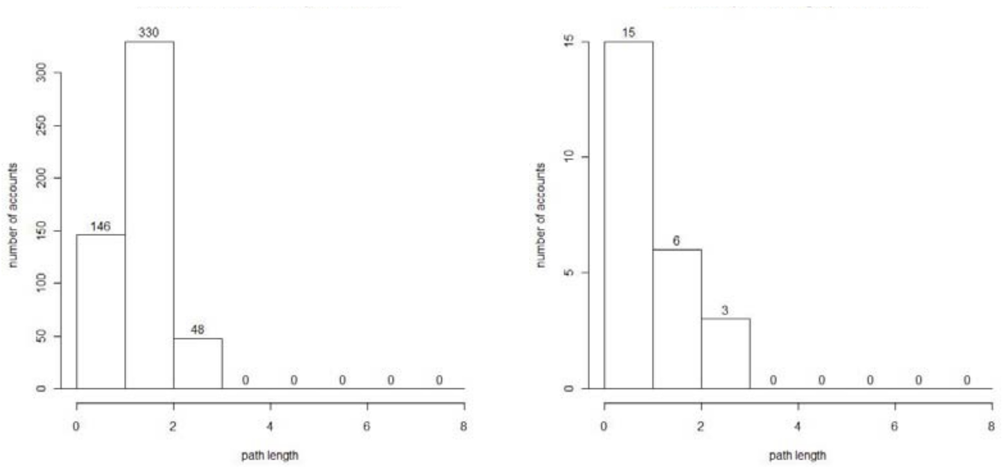

(a) normal accounts to fraudsters

(b) between fraudsters

Figure 2.1: Histograms of the shortest paths in the graph of Tunisia

each share was created. The observed techniques will be discussed in detail together with the detection algorithms below. On a high level the emloyed techniques involved at least two accounts, where one account traded shares for unreasonable prices with other accounts and all belonged to the same user. The minimum number of involved accounts was obviously two. The shortest path to the nearest fraudster was measured for each node using a breadth-first search. Two sets were defined: distance fraudster to fraudster and distance legitimate to fraudster. Since the underlying distribution is unknown but probably similar and the two sets are independent, the Mann-Whitney U test has been used to test the following hypothesis:

Hypothesis 2.2.1. H_0: The length of the shortest path between fraudsters is greater or equal than the shortest path between legitimate accounts and fraudsters.

Hypothesis 2.2.2. H_1: (Alternative) The length of the shortest path between fraudsters is shorter than the shortest path between legitimate accounts and fraudsters.

The p-values of the tests are listed in Table 2.1. The hypothesis obviously has to be denied for several shares at a highly significant level e.g. Australia, Ghana, Japan Korea, Mexico, Poland, Saudi Arabia and USA ($\alpha < 0.05$). But remembering the case presented in the introduction (Chapter 1 on page 1), there had been hard evidence of fraud in the Tunisia market. However, the Tunisia market graph seems to have other than the expected characteristics. With a p-value of 0.347 the hypothesis cannot be rejected for Tunisia.

Another surprise are the high p-values for the market shares of Angola, Argentina, Togo, Spain and Netherlands, where several cases of fraud had been found. There are several possible explanations for this: Either the finding of Cortes et al. does not hold for every market, or the fraudsters did only legitimate trades in these markets, or not all fraudsters in the selected network have been discovered yet. The latter explanation can be supported by the histograms of the observed shortest path distribution. Continuing the example

Table 2.1: Results of the Mann-Whitney U test on H_0 for each FIFA tournament share

Market	p-value
Angola	0.751
Argentina	0.775
Australia	0.049
Brazil	0.707
Costa Rica	0.900
Ivory Coast	0.388
Croatia	0.374
Czech Republic	0.723
Ecuador	0.529
England	0.155
France	0.148
Germany	0.180
Ghana	0.004
Iran	0.051
Italy	0.748
Japan	0.014
Korea	0.001
Saudi Arabia	0.024
Mexico	0.017
Netherlands	0.996
Paraguay	0.264
Poland	0.023
Portugal	0.105
Serbia & Montenegro	0.296
Spain	1.000
Sweden	0.249
Switzerland	0.993
Togo	0.798
Trinidad and Tobago	0.061
Tunesia	0.347
Ukraine	0.906
USA	0.026

of Tunisia, the corresponding histograms are shown in Figure 2.1. The vast majority of traders did not get directly in touch with a fraudster and have a shortest path length of two and above (c.f. Figure 2.1 (a)). On the other side, nine fraudsters seem to be separated from the other fraudsters at least by one or even two legitimate accounts (c.f. Figure 2.1 (b)). Since it is impossible to collude with only a single account, either these nine accounts had been labelled fraudulent due to activities in another share but for some reason "behaved" in this market or their partners have not yet been discovered. When analyzing the nine fraud cases, it turns out that all of them acted as legitimate accounts within the Tunisia market. Eight of the nine accounts only sold their initial endowment and then left the market. The ninth account showed only a little more activity by buying almost 3000 stocks during one week before the first match of Tunisia had started. But the strategies applied by the five convicted accounts appeared more often within the market - this time involving accounts labelled as legitimate so far. When changing the label of each newly discovered account the shortest path distribution shifted, bringing fraudsters closer together.

This is another example of how manual fraud detection with a small team on large data sets is an error prone and time consuming process: error prone because not all fraudsters had been found and time consuming because fraudsters have been discovered months after the markets had been closed. Therefore it is highly desirable to automate the process and minimize the necessary human involvement. Cortes et al. observed that fraudsters call friends and after they revealed to them how they bypassed the rules, these -so far regular accounts- then turned into fraudsters as well. On STOCCER, however, there is no such conversion. Multiple accounts have been created by the same person. Some of them were created to be able to play different strategies at the same time, while others were used to dump unwanted or worthless shares. In the transaction graph these accounts are naturally close together. The transaction graph contains important information about fraudsters and their relationships. So the open question is:

Problem 2.1. How can the market graph be used to detect fraud on prediction markets like STOCCER faster and with higher accuracy than human investigators?

The general idea is to mine this graph for unusual sub-graph patterns, especially patterns known from previous fraud cases. Each of these patterns is addressed by a particular heuristic optionally followed by a more detailed analysis in case of activation by the heuristic. The heuristics are used to enable the detection to happen in real-time or near real-time. The market graph is continuously updated with every contract concluded. The heuristics are started on each update. If the heuristic activates, it runs a detailed analysis. If the analysis detects an attempt at fraud, the case will be reported to the market operator. Each case consists of a score, the inflicted parties, the share, number of shares, transactions and volume.

Naturally, each indicator should capture a maximum variety of fraud cases and at the same time be as accurate as possible in its distinction between fraud and legitimate. In the following, two algorithms are going to be presented to detect two out of three observed patterns. The third pattern needs a different graph construction and will be covered briefly at the end of the chapter.

2.3 Ping-Pong Trading

2.3.1 Motivation

A common goal of market manipulators is to purposely push the prices once they have bought a sufficient amount of shares. One particular strategy is called *pump 'n dump* on stock exchanges. Via repeated selling and buying back the price of the stock rises. A high investment is necessary to buy existing orders in the order book. The hope is that the market starts to believe in previously undisclosed information / insider knowledge and jumps on the train. When enough people start buying, the fraudsters switch from buyer to seller and sell their stocks. Thereby they make profits from selling the stocks at the now increased stock price.

On STOCCER a similar pattern was observed but with a different goal. One account sold shares for a low price to buy them immediately back for a higher price within the spread. Consequently, the first account ruined himself step by step in favor of the second account. According to the resulting pattern the strategy was entitled *ping-pong* or *circular trading*.

Depending on the spread size and the current price of the share, this loop had to be repeated more or less often in order to transfer the initial endowment of one account to another. If the stock price is relatively high compared to the initial endowment, only few stocks can be bought and the process has to be repeated more often. In case of cheap or worthless shares, large quantities can be easily accumulated and only few trades are necessary. An example transaction sequence for such large scale transfers is given in Table 2.2 on the following page. Two accounts transfer ca. € 114.000 in less than 20 minutes with a *ping-pong* sequence of buy (b) and sell (s) transactions. Besides the alternating buy (b) - sell (s) sequence, the short time span between the two orders, listed in the last column, is very typical. The short time difference is necessary to prevent other traders from taking the orders. Though perfectly clean and almost synthetic, this case was found in the real world data set from STOCCER. Other instances are more difficult to spot since independent traders, who may have observed the activity in the order book, pushed orders in between and disturbed the pattern.

2.3.2 An Indicator for Ping-Pong Trading

As described above, the detection is split into two parts: a graph based heuristic and an in-depth analysis. The directed graph is built by representing accounts as nodes. Each

Table 2.2: Typical *ping-pong* transaction sequence for a cheap share

s/b	share	price	amount	volume	time-stamp	order.t diff.
s	Paraguay	€ 0.01	50518	€ 505.18	2006-06-16 09:43:28	00:01:12
b	Paraguay	€ 1.00	50518	€ 50,518.00	2006-06-16 09:46:18	00:01:38
s	Paraguay	€ 0.01	50518	€ 505.18	2006-06-16 09:50:30	00:08:14
b	Paraguay	€ 1.00	50518	€ 50,518.00	2006-06-16 09:53:37	00:01:26
s	Paraguay	€ 0.01	50000	€ 500.00	2006-06-16 09:55:56	00:13:40
b	Paraguay	€ 0.29	50100	€ 14,529.00	2006-06-16 10:01:12	00:03:12
	buy-sum			**115,565.00**		
	sell-sum			**1,510.36**		
	profit			**-114,054.64**		

edge from a arbitrary node a to node b represents all transactions so far concluded where account a has sold to account b. For a real-time system, the heuristic is called for each updated or newly created edge. For an ex-post analysis, it is just called once.

Ping-Pong Trading Heuristic

This indicator detects cycles in the graphs. In the basic implementation, it only detects cycles of length two which involve just two nodes. The reason this simplification was desired is the real-time aspect of the heuristic. Besides, the effort for the fraudster to implement a strategy with more than two accounts is much higher and thus less likely. To improve the speed of the heuristic, a separate graph for every traded stock is created. This assures that whenever a cycle is detected, the heuristic does not have to test whether stocks of the same share have been traded. This particular fraud technique does not work across various shares as after the first sell, all shares would be gone and the user would not have enough resources to perform a second cycle. Large isolated deals in favor of another account will be topic of the next presented indicator below.

Another advantage of separate graphs per stock is that a case will only be reported once to the user for every stock-seller-buyer combination — even if the *ping-pong* strategy has been used several times. While running the market and updating the graph, two possible cases can occur:

1. A new edge from node a to node b is inserted into the graph.
2. An existing edge of the graph from node a to node b is updated.

In both cases, the heuristic checks whether there exists an edge from b to a which closes the cycle. If this is the case, and the last modification time-stamp difference between the two edges is lower than a certain limit ($\max_{\Delta t}$), a detailed check of the transaction history is invoked. The introduction of the time limit is required to revise existing edges as well. This prevents the heuristic from detecting cycles based on trasactions which are two long ago and of rather speculative nature where the market developed into another direction than assumed by the trader. Switching from buying stock back to selling stocks is natural in this case and not desired to be detected by the heuristic. This is covered by case 2.

In-depth analysis

One could imagine that quite a number of possible tests are necessary to classify a case as *ping-pong* trading. For example, looking at the number of sell and buy transactions, the number of stocks exchanged between two traders, the time difference between buy and sell transactions, price differences and so on. To keep the test as simple as possible, the transaction history is filtered for the currently traded stock and gets ordered by the execution time. Stepping through the history, the average price for each block of sell transactions and for each block of buy transactions is calculated. Whenever the trading direction switches from sell to buy or vice verse, two checks are performed:

1. Is the time span between the last buy/sell and the current sell/buy transaction lower than $\max_{\Delta t}$?
2. Is the average sell price lower than the average buy price? In other words, is the price on edge $\overline{ab}$ lower than on the edge $\overline{ba}$?

If both conditions are fulfilled, a score l is incremented. The score correlates to the probability that the underlying account pair are fraudsters. The intuition of this in depth analysis is to count unprofitable changes of trading direction which have a high probability of not being caused by changing long term investment plans. The latter is addressed by the first condition stating that the transactions have to be concluded within a short period of time. The second condition also assures that only one case per cycle is reported. The current implementation puts the emphasis on the one who is ruining himself by looking for unprofitable transaction sequences. Both accounts are reported per case. A case is only reported when the incremented score is higher than a previously selected level L. How often someone may change their mind or conduct a transaction due to an input error and tries immediately to correct it, depends on the underlying market and trading interface.

Analysis of complexity

The indicator, comprising of the heuristic and in-depth analysis, is heavily depending on the number of nodes compared to the number of transactions. The heuristic will only fire in case a circle is closed. The probability of a new inserted edge cloosing a loop increases with the density of the graph $G = (V, E)$. The density is the proportion of edges relative to the total number possible $|E|/(|V| * (|V| - 1))$. For the online sports prediction markets like STOCCER with open registration the major part of the graph is very sparse and there is only a small core cluster of frequent traders. Furthermore, sports

Table 2.3: Number of reported cases varying $\max_{\Delta t}$ and setting $L = 0$

$\max_{\Delta t}$	# reported cases
∞	246
240min	134
120min	127
105min	127
90min	127
75min	125
60min	119
45min	117
30min	111
15min	93
10min	80
5min	74
1min	41

prediction markets are operated only for a limited time, such as till the end of the league, championship or cup. This means that the density may vary but is likely to increase towards the end. Across the 32 different shares traded on STOCCER the average final graph density was 0.00933.

The in-depth analysis is more complex but can be further optimized for online processing via only going one buy/sell block back in history as the older ones have been checked in previous alerts of the heuristic. In case the previous transaction was in the same direction we have already performed a similar check and we can abandon. In case our database storage returns the transaction in their natural order which is sorted by time, the complexity can be seen as constant ($O(1)$) since the number of transactions to check is very limited.

Summarizing the above, the complexity of this indicator is $O(\frac{|E|}{|V|^2})$ for an online version. In the offline ex-post analysis we add an important constant which is the average number of transactions per edge. In sparse markets like stoccer this number is often small because of the many one day visitors who only have few trades and lower the average.

2.3.3 Applying the Indicator

The optimal fine-tuning of the indicator requires a sensitivity analysis. The indicator has two parameters: the maximum time between a change in the trading direction ($\max_{\Delta t}$) and the minimum repetition threshold (L) above which a case is reported. Table 2.3 lists the number of reported cases in the STOCCER data set when $\max_{\Delta t}$ is varied while L is kept fixed at 0. If no time restrictions are set at all, 246 cases are reported. This means in 246 occasions someone had made an unprofitable sell-buy- or buy-sell-combination with the same trading partner.

Table 2.4: Number of reported cases varying L while setting $\max_{\Delta t} = \infty$

L	# reported cases
0	246
1	97
2	56
3	42
4	26

Since STOCCER was about predicting the outcome of soccer matches, the question arises how fast people adapted to unexpected outcomes of the matches. If people would trade their opinion before the match and return after watching the match, a significant change in the number of reported cases is expected at around 105 minutes (2 * 45 minutes plus 15 minutes break) or at least 45 minutes in case they are trading during the break. This is not the case. The number of cases reported stays high and only drops significantly if $\max_{\Delta t} \leq 10$ minutes. This indicates that most of the participants traded while following the match. From a prediction market perspective, this makes sense as newer information is getting available during the match (goals, fouls, substitutions, etc.). It is surprising that the high number of 41 cases remain even after the maximum time-span has been decreased to one minute. A further reduction of the time-span below one minute has not been made since on a web-based trading interface like STOCCER the time for the update of the order book plus the time of filling out the order form and committing it can sum up to more than half a minute. This depends on the overall load of the servers and bandwidth. One explanation could be that the user prepared the two transactions in two different browser windows to submit the orders as quickly as possible.

In a second step, the influence of the minimum repetition threshold L was tested, while $\max_{\Delta t}$ was set to infinity. In STOCCER, between 500 and 900 accounts traded a share. That one account traded twice with the same trading partner is unlikely. As expected, L turned out to be an effective filter (see Table 2.4). Only 56 cases had more than 2 unprofitable cycles and 26 had more than 4. Remember, a cycle means that an account sold cheaper to a second account than the price for which it bought the shares back from that same account. Thus, a L-value of 4 means that the same two accounts (i.e. a, b) had 4 of these unprofitable cycles always in favor of the same account (w.l.o.g. a). The highest observed repetition has been observed for Tunisia with $L = 26$, indicating that two accounts shifted money from one to the other in 26 cycles . For $L \geqq 27$ no case are reported anymore.

2.3.4 Choosing Optimal Parameters

Each market is different. The analysis above indicates that depending on the number of traders, the trading frequency, and the underlying, the parameters L and $\max_{\Delta t}$ have to be adjusted. A general recommendation for the two parameter settings cannot be given

Table 2.5: Example transaction sequence of the account *willi* for a selling cheap

to	b/s	share	price	amount	volume	time-stamp	t. diff.
Azzura	b	France	10.48	300	3,144.00	11.06.2006 01:01	01:01:59
Hainz	b	France	10.55	200	2,110.00	11.06.2006 01:01	01:13:28
Baschi	b	France	10.9	6626	72,223.40	11.06.2006 01:01	00:01:25
waba554	b	France	10.13	374	3,788.62	11.06.2006 01:01	00:36:38
Azzura	b	France	10.33	300	3,099.00	11.06.2006 01:01	01:02:11
Baschi	s	France	8.6	7800	67,080.00	11.06.2006 01:08	00:01:17
		buy-sum			**84,365.02**		
		sell-sum			**67,080.00**		
		profit			**-17,285.02**		

and as the small example of the expected trading time of 105 minutes indicates, a sound knowledge of the market is necessary to separate the wheat from the chaff. At the end of the chapter a combined approach of the two indicators is presented where no restrictions have to be used at all.

2.4 Prominent Edges

2.4.1 Motivation

Ping-pong trading leaves quite obvious traces in the transaction history. Besides the repeated money-pumping addressed by the indicator above, a second strategy was frequently observed as well. The second strategy tries to cover the involved accounts by reducing the number of necessary transactions to a minimum. The idea is to get rid of worthless stocks for a good price or to buy expensive stocks far below the average market price. Consequently the counter party is ruining itself. An example of selling worthless shares has been presented in the introduction in the previous chapter. Table 2.5 illustrates a case of selling cheap with a transaction listing from the STOCCER data set. The last column *t.diff* indicates for how long the matching orders from the account mentioned in the first column have been in the order book. The participant *willi* is buying a high number of France-stocks with a single large order, indicated by the same time-stamp of the transactions, then selling them afterwards to *Baschi* far below the price he paid just a few minutes ago. *Baschi* is gaining 7800 shares of France for 17,285 less than the current market price. This happened several days before the first match where France played for the first time. It is very unlikely that new information had been revealed at that very moment, especially so late at night.

The pattern of this strategy in the graph is an edge with a significantly higher volume than other edges of the same node, since the supporting account tries to give the maximum profit to the selected account. In the case of transferring stocks below market prices, the

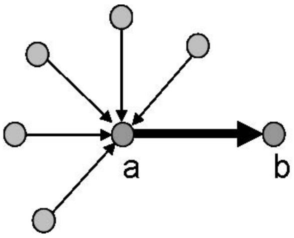

Figure 2.2: Prominent-Edge graph pattern

expensive stock has to be bought first by the account that is selling the stocks afterwards. This results in a star pattern like the one depicted in Figure 2.2.

2.4.2 The Prominent-Edge-Indicator

The challenge is to identify trading volumes deviant from the regular volume traded by the account. The difficulty is that every trader has a different trading strategy and in some cases, also different financial resources. Moreover, different trading volumes may be induced by the platform if it provides different initial endowments on each market. Therefore, the threshold to classify an edge as abnormal has to be adapted for each market or even each node. To overcome this problem, we employ a concentration measurement. Concentration measurements can be categorized into relative (e.g. Gini Coefficient (Gini 1936)) and absolute measurements (e.g. Herfindahl). While relative measurements only compare each edge with the others using a selected attribute, absolute concentration measurements focus more on the number of edges.

This distinction is especially relevant for a leaf: a node with just one single edge. A relative measurement assigns a value of zero to the edge, while evaluating the leaf node, because there is no other edge available for comparison. An absolute measurement assigns a value of 1.0 to the node, because the complete volume is concentrated on only one edge.

Online platforms often have a high number of one-time users. These are users who just want to explore a new service or platform and loose interest shortly after their first login. This also holds for STOCCER. Many users only traded during one day and then left the market without returning ever again. Figure 2.3 depicts on the x axis the number of days of activity. From the 444 one-day-users 70 had only one single transaction turning them into leaves. If an absolute concentration measurement is applied, it would make these 70 one-time users highly suspicious, hence they receive the highest possible score of 1.0. One could think of special strategies or treatments of these leaf accounts, such as testing them with a special heuristic. Instead, a relative concentration index is used since it outperforms the absolute concentration measurements. The latter ones are less sensitive to relative size differences and more to the absolute number of edges. A fraudster could use this omission for his profit by creating many accounts all of which are used only for

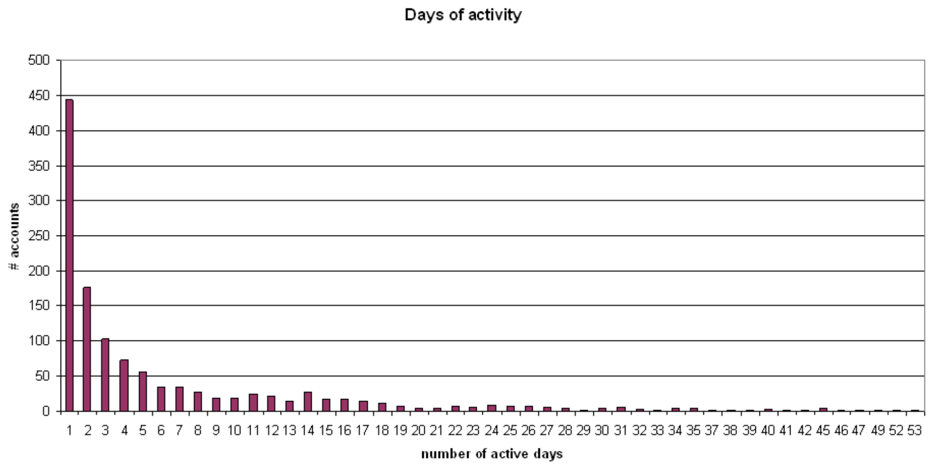

Figure 2.3: Histogram of the distribution of the number active days per accounts in the STOCCER data set

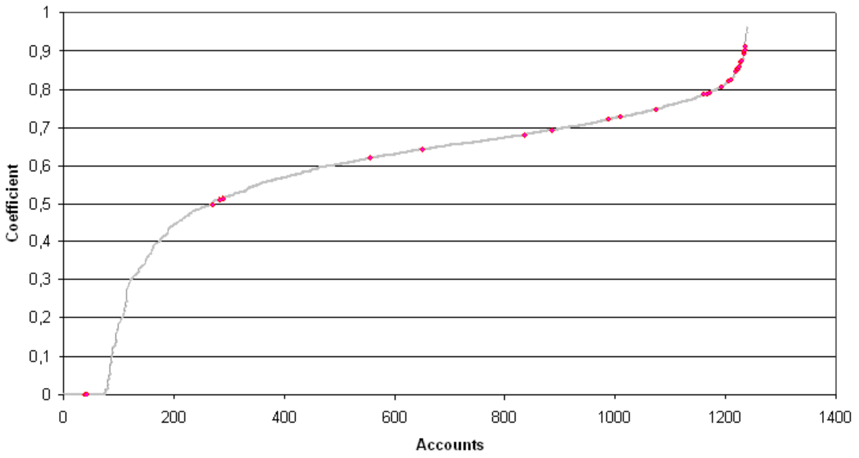

Figure 2.4: Gini relative concentration measurements on the volume in the STOCCER data set

a single transaction in which he buys/sells to his primary account. However, the pattern of an account with an unusually high number of leafs, again makes this very obvious in the graph and easy to detect (e.g. via a breadth first search).

The result when calculating the Gini Coefficient based on the traded volume per edge without any parameterization is displayed in Figure 2.4. Along the x-axis the accounts are sorted by their coefficient value. The red dots indicate that the corresponding account had been found during market operation and tagged as fraudster. It is obvious that besides some outliers the majority of fraud cases have a coefficient value higher than 0.75. For a coefficient greater than 0.75 there are 145 cases reported; for a coefficient greater than 0.8 there are only 52 cases. Of those, 13 out of the 52 cases were previously classified as fraud. While revising the remaining 39 accounts, 11 of them could be easily classified as fraud. But still more than half (54%) are false positives.

These false positive accounts had issued an order which was matched days or even weeks later. This is another consequence of one-time users, who place orders which may stay a long time in the order books. Somebody may have put a sell order for a low price long before the first match. If the team later unexpectedly wins and the prices rise within minutes, the order may appear uninformed and gets cleared. This can create very strong edges between unrelated accounts. Such occurrences can be seen more as a speculative strategy and should not be addressed. These edges of unrelated trades have to be filtered before applying the Gini coefficient. Therefore we construct an undirected graph using only transactions where the matching orders arrived within a certain time span, e.g. time-stamp of order A - time-stamp of order $B \leq 1$ hour.

The full indicator can be described as follows:

1. Filter all transactions/edges where the timestamps of the corresponding buy and sell order differ not more than $pe.\max_{\Delta t}$.
2. Convert directed edges into undirected edges. In case of an edge $e_1(a, b)$ and a second edge $e_2(b, a)$ remove both and introduce an new undirected edge e_3 where all attributes are the sum of the attributes of e_1 and e_2.
3. Calculate for the filtered edge-set the Gini coefficient:

$$G(v) = \sum_{e_i \in edges(v)} \frac{2i - n - 1}{n} * e_i.attr$$

   with v being the node currently evaluated, and $attr$ the selected attribute the concentration is based on. Depending on the market and the distribution of prices $attr$ may be the number of transactions, the number of shares or the volume ($\sum price * quantity$) traded over the edge.

If $G(v)$ exceeds a threshold θ, the case is reported to the market authorities. The conversion of the graph into an undirected graph helps to distinguish cases where a node traded on several edges with similar volumes, but on one edge received a high volume in return. In the directed case, this would not call the attention of the authorities though certainly being worth a look.

The runtime complexity of this indicator depends on the implementation being either focused on realtime or ex-post analysis. While for a real-time analysis the graph update only happens in case the current transaction matches the filter criteria, the ex-post analysis will have to build the full graph for all transactions once in the beginning. The complexity of the Gini coefficient is per evaluated node equal to the degree of the node. The upper boundary is naturally $O(|V|)$ for a node with an edge to every other node in the graph. In STOCCER the maximum observed outdegree was 200 for the Iran market with 614 nodes. This is an extreme exception as the majority of nodes had less than 20 edges (see Figure 2.5). Given the logarithmic degree distribution in such markets the complexity is between $o(\log |V|)$ and $O(|V|)$. For an expost analysis the whole graph will

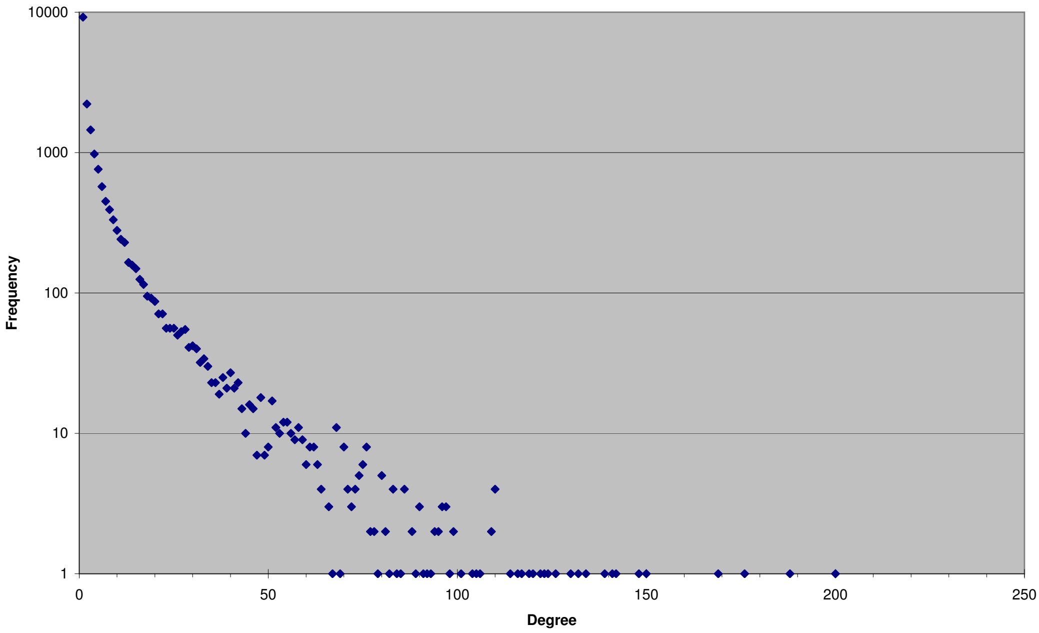

Figure 2.5: Histogram of node degrees

be traversed, which adds the total number of nodes as a multiplier for the upper bound $O(|V|^2)$ and expected complexity $o(|V| \log |V|)$.

2.4.3 Applying the Indicator

Table 2.6 shows the sensitivity of the parameter $pe.\max_{\Delta t}$. Obviously the number of reported cases can be reduced significantly. The unexpected increase in reported accounts while reducing the time span towards 12 hours is related to the graph construction: Some accounts had several strong edges. But when we reduce the time span, some transactions are removed from the transaction history and at a certain span only one strong edge remains in the transaction history of the account. Now this account will be reported until the time span drops below the value of the last strong edge.

The unexpected increase underlines that indicators probably fit best if they are calibrated node-wise and not just according to the market. Certainly an optimal calibration cannot be achieved manually, since this would be an immense effort requiring a complete manual investigation of the whole data set.

Since many of the false positives resulted from different sell and buy transactions in different shares, some modifications have been implemented and tested. The following modifications have been implemented:

1. a directed graph,

Table 2.6: Number of reported accounts varying $pe.\max_{\Delta t}$ and keeping $\theta = 0.75$

$pe.\max_{\Delta t}$	# accounts
24h	145
22h	139
20h	138
18h	138
16h	138
15h	140
14h	147
13h	141
12h	142
10h	137
08h	131
04h	119
03h	122
02h	108
01h	87
45m	83
30m	72
15m	59
10m	49
05m	44
01m	16

a) considering the entire graph

b) considering only incoming edges

c) considering only outgoing edges

2. setting up different graphs for each share.

As expected the directed graph implementations returned fewer cases than the undirected version. Unfortunately, there was no notable change in the false-positive rate (more or less 50%). The implementation using a different graph for each share got the best result with only 32.5 % false positives ($pe.\max_{\Delta t} = 1h,\ \theta = 0.80$). The logic behind this effect can be explained by the lower variance of the prices. If an account sells his complete initial endowment to the market, he makes less profit for cheap shares and a higher profit for the more expensive ones. If only transactions of the same share are considered, the concentration is related to the amount of stocks. If an account usually trades only less than 500 shares and suddenly 5000 with a single partner, the indicator produces the desired warning. The resulting case list is provided in Table 2.7. The table contains the case list with the previously described settings. The accounts are colored according their manual classification into fraudsters (red) and legitimate (black) accounts. The table is ordered according to the edge volume instead of the Gini-value (column score). The column labelled *tr.* lists the number of transactions and *q.* refers to the number of

Table 2.7: Cases reported by the Prominent Edge Indicator per share $(pe.\max_{\Delta t} = 1h,\ \theta = 0.80)$
Note: Fraudsters are marked in red color.

account 1	account 2	Gini Indicator for	score	*volume*▽	tr.	q.
FEZ	RibAldA	[Côte d Ivoire]	0.8294	540.0	1	200
FEZ	RibAldA	[Czech Republic]	0.8060	1900.0	1	200
gunnar.kaestle	Maddin	[Serbia & Montenegro]	0.8039	3348.0	3	1200
Soccer securities	Studienrat	[Togo]	0.8074	9719.35	3	8673
bleckfriday	Bernie78	[Czech Republic]	0.8650	10,000.00	1	1000
brainjohn	hannes2802x	[Costa Rica]	0.8030	12,979.00	2	12595
KöbiKuhn	Andrea	[Costa Rica]	0.8275	13,500.00	1	7500
FEZ	sb3000	[Germany]	0.9087	15,815.15	2	815
HelloWorld	eug555kg	[Poland]	0.8337	21,930.00	1	12900
winner	Hopp Schwiiz	[Trinidad &Tobago]	0.8150	23,270.00	1	13000
Doppelpack	KSC4ever	[Angola]	0.8564	23,350.00	1	23350
brainjohn	bleckfriday	[Iran]	0.8041	23,583.00	1	23583
JPKocher	dagho	[Ghana]	0.8080	26,590.50	1	5598
welwelsken	merlin	[Togo]	0.8191	27,531.00	4	18600
rizzopower	king	[Paraguay]	0.8506	35,340.74	3	93002
bleckfriday	Liceu	[Iran]	0.8817	37,568.21	6	23981
KSC4ever	blablu	[Côte d Ivoire]	0.8159	39,996.00	1	9999
www.tischt. . .	Attila	[Angola]	0.8418	45,843.43	4	48758
RibAldA	pabeki	[Paraguay]	0.8199	46,136.40	2	11504
drogadito	morros	[Tunisia]	0.8831	47,723.01	4	23909
www.tischt. . .	Attila	[Korea Republic]	0.8333	48,018.81	1	19131
fpschebe	falke	[Ghana]	0.8595	54,320.00	2	21728
RibAldA	Mikl	[Trinidad & Tobago]	0.8836	66,570.00	1	4438
Maddin	Pia	[Angola]	0.8461	81,155.84	5	37532
fruit	king	[Paraguay]	0.8624	114,000.00	2	120000
slindoe	kyrie	[Tunisia]	0.8042	213,014.94	17	97639
kyrie	slindoe	[Korea Republic]	0.8134	326,750.62	28	73082
Maddin	Lutscher	[Ghana]	0.8476	585,409.00	25	118050

stocks transferred between the counter parties. Obviously, the higher the edge volume is, the higher the probability of fraud. This strategy is very convenient for a fraudster. The fraudster has to create fewer fake identities or cooperate with fewer partners and has to transfer less often but in higher amounts to improve the chances of the primary account. On the downside, this increases the probability of false positives for cases with a low volume.

This finding has been reproduced on further data sets. Table 2.8 lists the results of the *Prominent Edge Indicator* applied to four data sets. The parameters were kept fixed with only the FIFA WM 2006 market being evaluated twice. Once with the default setting applied to all data sets and a second time with the example settings above to make it more comprehensible in combination with Table 2.7. The volume median $(\widetilde{x})$ for Table 2.7 is 31,435.87 and the 0.25 quartile $(x_{0.25})$ equals 15,236.36. In all tested data sets, the

Table 2.8: Reduction of fraudsters vs. false positives (FP) by cutting of the low volume cases at 0.25 quartile or median

$pe.\max_{\Delta t} = 0.75,\ \theta = 24h$	**all cases**		**cut $x_{0.25}$**		**cut $\tilde{x}$**	
data set	**#fraud**	**#FP**	**#fraud**	**#FP**	**#fraud**	**#FP**
FIFA WM 2006	61	108	60	77	56	45
1. Bundesliga 06/07	3	4	3	1	3	0
1. Bundesliga 05/06	8	10	8	7	8	2
Herbstmeister	3	6	3	6	3	5
$pe.\max_{\Delta t} = 0.80,\ \theta = 1h$						
FIFA WM 2006	27	13	23	10	18	3

number of detected fraudsters stays fixed or is only slightly lower. However, the number of false positives can be significantly decreased.

2.5 Combined Approach

Although the reduction in reported accounts by parameterization is very successful for both indicators, it obligates the market supervisor to carefully evaluate and set the parameters. Without any parameter restrictions ($L = 0$ and $\max_{\Delta t} = \infty$) the *Ping-Pong Indicator* marks 246 account pairs as suspicious. These suspicious cases can be fraud or simply the result of a change in the trading strategy or an input error. It is desirable that a pre-selection of the parameters – similar to the median cut-off in the *Prominent Edge Indicator* – could be applied, which filters the highly suspicious cases. A ruinous buy-sell combination is suspicious, especially when it is one of the biggest transactions in terms of volume for at least one of the accounts.

This suggests combining the two indicators. In a first step, both indicators are run independently. Secondly, all edges reported by only one indicator are removed from the result set. Even with very open parameter settings, amazing results were obtained. The *Prominent Edge Indicator* time span was set to 12h, taking only accounts with a coefficient greater than 0,75 into consideration. The *Ping-Pong Indicator* was set to $L = 0,\ \max_{\Delta t} = \infty$, reporting 246 cases. After removing all cases with a unique account combination – meaning that a specific account combination was reported only by one indicator and only once –, 52 accounts remained in the list. Of these, 19 of them had been found already during market operation. A hand revision of the 33 remaining accounts confirmed the suspicion in 28 cases, but was indeceicive with regard to the last five. It was unclear if the traders made this transaction sequence by accident or on purpose. They were reported because of having an unprofitable buy-sell combination in two different shares to the same counter party.

Nevertheless, it is a valuable option to avoid the evaluation of parameter settings and still have an effective profiteering, reducing the false positive rates below 10% (compared to 50% of the pure prominent edge).

2.6 Detection of Robots/Scripts

Besides the two strategies addressed in the previous chapters, the terms and conditions prohibited specifically the usage of robots or scripts by the participants. The goal was to have a manual trading market. Some cases however provide evidence of script usage. For example, one user was trading at a very high frequency during several hours. The average inter-arrival time of his orders was lower than 10 seconds. His orders covered all 32 markets and -particularly suspicious- in alphabetical order, one order per market and all orders of the same lot size. The lot size changed after each alphabetical iteration. Similar patterns have been observed with other accounts. Though the user interface would allow to sort the shares in an arbitrary order, the manual trading of the same volumina in all shares does not make much sense as portfolio trading was offert via the trading interface as well.

In order to detect these cases automatically, a time series analysis of the transactions is proposed. A graph based solution is to create a bipartite graph of accounts and shares. When an account trades the same set of shares within the same amount of time, a star like pattern arises. If the pattern repeats, the operator gets an alert.

However, every fraudster reading this section will modify his robot in such way that it will not be detected by the above suggested methods. A more general approach to exclude scripts from trading is the usage of CAPTCHAs (Ahn et al. 2004a,b). The user has to solve a hard artificial intelligence problem when logging into the system. Since the human abilities in various areas are still unmatched by the machines, one of these cases can be used to exclude scripts from logging into the system.

2.7 Further Possible Extensions

2.7.1 IP-based Approaches

So far, only transaction data have been used for the indicators and their validation. Another source of information is the IP-address submitting the order. This information may put further weight into the scale. Maranzato et al. used workstation identifiers and ips in combination with certain correlated market events (Listing + Buyer) two defraud reputation systems for an electronic market in Latin America (2009; 2010). Many fraudsters may ignore the fact that the network connection is giving another trace. However, in IPv4 networks this may not be a unique identifier due to gateways, proxies, and VPNs. Careful fraudsters will always use two different IP-Address to commit fraud. Therefore this information has been ignored in the present state of the indicators.

Table 2.9: Example creation pattern of a collusion group

id	account	name	surname	email	address	cc
1603	stoccer	bbbbbbbbbbb	blöd	schweizwollerau@***.net	ffffffffff 4455 zuu	17
1594	wollerau	***[1]	****[1]	abcde1@***.ch	hhhhh 33 8800 Zürich	188
1597	123wollerau	bbbbbbb	bbbbbbbbbb	123wollerau@***.ch	ggggggg 8800 z	17
1595	wollerau123	bb	cccc	wollerau123@***.ch	hhhhhhhhhhh 8800 Zürich	28
1601	123tim	hhhhhhhhhhh	hhhhhhhh	wollerauschweiz@***.ch	hhhhhhhhhhhhhhhh 8880 zuri	17

[1] characters replaced by * for anonymity

2.7.2 De-duplication on User Registration Data

Generally, one would expect fraudsters to try their best to create independent identities. Maranzato et al. found that even just the domain of the registration email might be a good signal in case of existing suspicion (2009). Surprisingly, on STOCCER many cases have been found with very similar registration data. Even if the data itself was not identical, at least the way it was altered was similar. Table 2.9 lists five accounts. The selected country was stored as an integer in the database under the abbreviation *cc* for country code. Obviously, the data used to register these account are highly similar. But especially the way they differ supports the hypothesis that they originate from the same person.

Unfortunately, it is still an open question to find this kind of pattern algorithmically. So far, only the human pattern recognition abilities can help to identify and use these patterns to strengthen the evidence against a suspicious group of accounts. Therefore, the implemented prototype (see chapter 4 on page 71) offers an incremental search over all fields of the registration data. This allows the user to efficiently check suspicious patterns.

2.7.3 Neighborhood Search

As described in the section 2.4.2, a relative concentration measurement like the Gini coefficient ignores leafs, because there are no other edges with which to relate. But fraudsters do exist within the group of leafs as Figure 2.4 indicates. For these special cases, the pre-filtering via the volume has been tested as well. The results are promising, with 19 of 82 edges stay in the list when cutting away all edges with a volume below $x_{0.75}$. Within the remaining set of edges, eight fraud accounts have been found. Six out of eight could be confirmed via similar registration data. The last two accounts have been found guilty because of convincing transaction patterns.

Regarding the false negatives this is still poor evidence, but all of them are partners with already detected fraudsters. Therefore, a neighborhood search feature has been integrated into the prototype, allowing it to filter the indicator result list for cases where at least one of the two accounts already is flagged as a fraudster (see 4.2.1 on page 79).

2.8 Summary

In this chapter two graph based approaches to detect collusion in prediction markets have been presented. The goal was to find simple measurements that are fast to employ and easy to control by the market authorities. The influence of the different parameters on the precision has been evaluated. To avoid the calibration of parameters, a combination of both indicators has been proposed. Thereby a precision of above 90% could be achieved on a real world data set.

Naturally, using graphs implies higher memory consumption. On the other side, a high run-time performance can be achieved. The *Circular Trading Indicator* needs less than seven minutes for a complete replay of all 32 graphs with a total of 17,432 nodes and 47,588 edges; the *Prominent Edge Indicator* even less than 30 seconds.

A side effect of the investigation was the finding that people trade upon each goal instead of betting on the outcome of the game in advance. The high number of ruinous trades within less than half an hour indicates that orders were issued alongside the goals of the matches. Consequently, some of the traders bet on the wrong horse and tried to cut their losses. This also implies that the most active traders watched the matches with their computers near by.

Cortes et al.'s findings (2001) about fraudsters being closer to other fraudsters than to normal users have been extended from telecommunications to markets. Besides the first test in the beginning of this chapter, a second test with an extended fraud account set has been conducted. This set included the accounts which were manually excluded during the market run-time, plus the newly discovered accounts via the combination of the *Circular Trading* and the *Prominent Edge Indicator* (see Section 2.5). Again the shortest path from each account to the next fraudulent one has been measured and the hypothesis H_0 (the length of the shortest path between fraudsters is greater or equal than the shortest path between legitimate accounts and fraudsters) tested using the Mann Whitney U Test. From the subset of fraudsters detected by the algorithms, only accounts detected in the corresponding share market were used. In other words, if an account had been found because of fraud in the Angola and Swiss market, the account was added to the manual detected set only for Angola and Switzerland, but not for other markets.

The result is depicted in the second data column in Table 2.10. In the manual classified set, only eight of the 32 markets had a p-value < 0.05 (see Section 2.2). After applying the indicators the hypothesis H_0 cannot be rejected for 21 markets. The p-value for the

Table 2.10: Results of the Mann-Whitney U test on H_0 for each FIFA tournament share

Market	p-value[1]	p-value[2]
Angola	0.751	<<0.000
Argentina	0.775	0.026
Australia	0.049	0.001
Brazil	0.707	0.053
Costa Rica	0.900	0.074
Ivory Coast	0.388	0.003
Croatia	0.374	0.048
Czech Republic	0.723	0.265
Ecuador	0.529	0.018
England	0.155	0.136
France	0.148	<<0.000
Germany	0.180	0.004
Ghana	0.004	<<0.000
Iran	0.051	0.029
Italy	0.748	0.083
Japan	0.014	0.005
Korea	0.001	<<0.000
Saudi Arabia	0.024	<<0.000
Mexico	0.017	<<0.000
Netherlands	0.996	0.055
Paraguay	0.264	0.005
Poland	0.023	<<0.000
Portugal	0.105	0.100
Serbia & Montenegro	0.296	<<0.000
Spain	1.000	0.079
Sweden	0.249	0.129
Switzerland	0.993	0.992
Togo	0.798	0.001
Trinidad and Tobago	0.061	<<0.000
Tunesia	0.347	0.003
Ukraine	0.906	0.199
USA	0.026	0.001

[1] p-value for the set of manually excluded fraudsters (c.f. Table 2.1)

[2] p-value for the set of manually + algorithmically excluded fraudsters

remaining 11 markets improved significantly with five being $<= 0.10$. Only for the Czech Republic, England, Sweden, Switzerland the hypothesis could not be rejected.

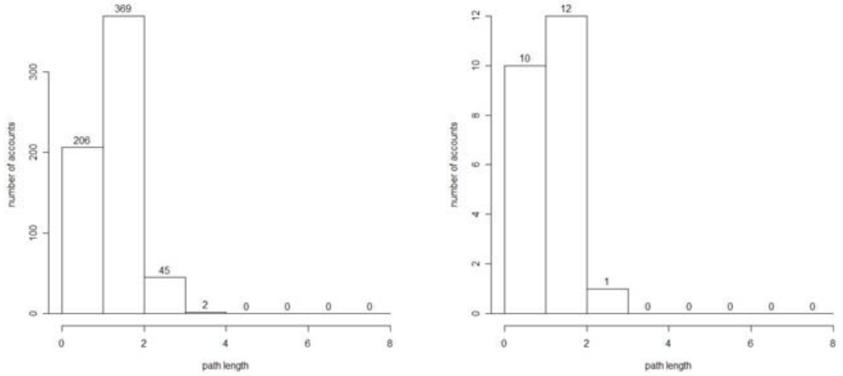

(a) normal accounts to fraudsters

(b) between fraudsters

Figure 2.6: Histograms of the shortest paths in the graph of England

Since in the second test the manual set was still used for all markets, the shortest path callculations might not be optimal since the underlying activity for some of the edges had been legal. When looking at the histograms for these five markets, the distributions look very similar for the shortest path from normal to fraudsters and from fraudsters to fraudsters (see as an example Figure 2.6 for the England market). Only 7.5% of the honest traders in the English market had a shortest path longer than two to the next fraudster (Figure 2.6 (b)). This indicates that many fraudsters played a central role as major traders on these markets. When visualizing the England market graph and marking all detected fraudsters in red, this suspicion is confirmed (see Figure 2.7). Many fraudsters are in the center of the graph, being the major trading partners for other fraudsters and honest traders. As a result, the shortest path for honest traders is rather short and the Mann Whitney U Test cannot see them beeing significantly different compared to fraudsters.

The question, whether this observation defeates the applicabillity of Cortes et al.'s findings, has to be answered with 'no'. Having fraudsters operate as major market players is a result of not monitoring market acitvity for possible fraud. This happened only since STOCCER was primarily a scientific experiement about prediction markets. The need for rule enforcement and market monitoring was only discovered during the experiment. But without proper tools there were still plenty of users who managed to abuse the system while staying under the radar. Some of the fraudsters are neighbors of the manually excluded fraudsters. Hence, the results of the mixed test show a definite improvement. This shows that Cortes et al.'s findings also hold true on the STOCCER platform and that the suggested indicators help to reveal fraudsters in an automated way.

Moreover, there is the possibility of still undiscovered false negatives, leaving room for the development of further indicators (e.g. robot detection) and the improvement of

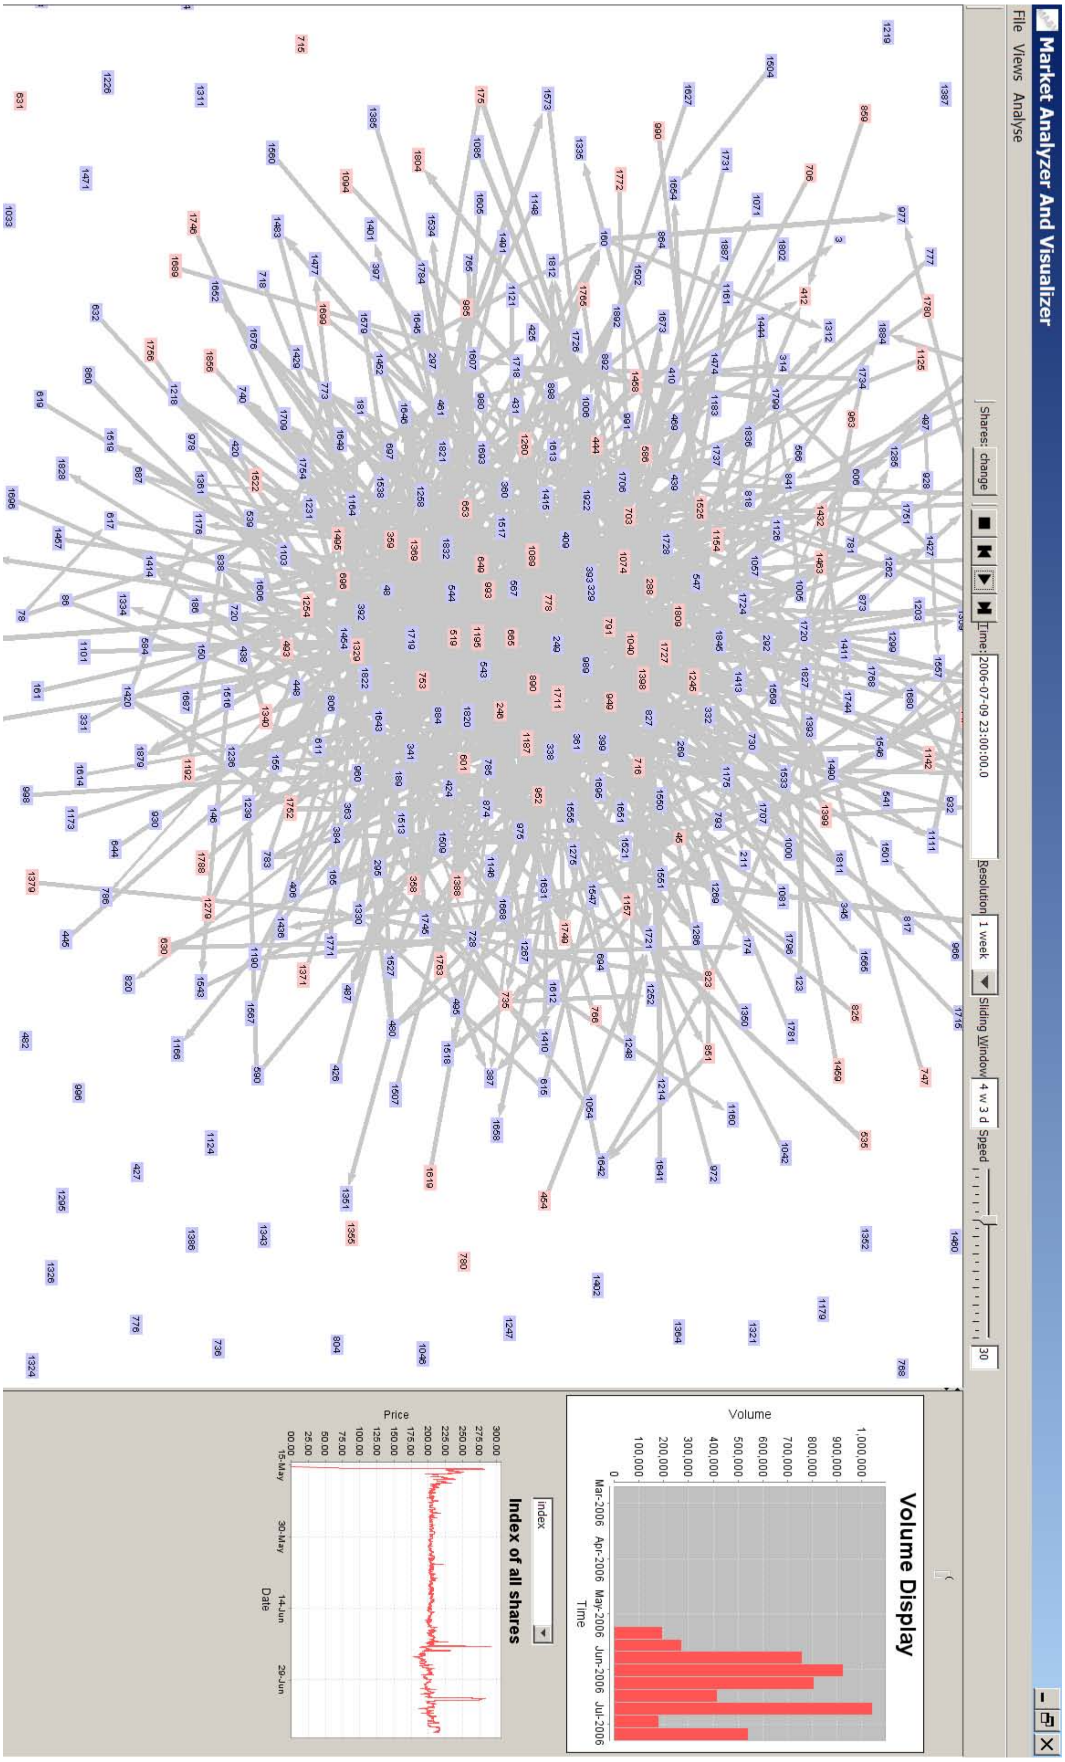

Shares: change

Time: 2006-07-09 23:00:00.0

Resolution 1 week

Sliding Window 4 w 3 d

Speed 30

Figure 2.7: Marketgraph for England of the last four trading weeks

existing ones as explained in 2.7. An completely independent evaluation of the algorithm results will be presented in Chapter 5.

3 Market Visualization Design

Markets are driven by price and liquidity. Therefore, common visual presentations of markets focus on price and volumes. Some modern visualization technologies such as treemaps (Turo and Jungmeister 1992) and information landscapes (Gershon and Eick 1997) have been applied to markets without becoming widely established. All these visualizations address the trader and his needs, but are insufficient for market operators and supervisors, since they do not provide deeper insight into the trade activity. For example, the question of how reliable the aggregated information is cannot be answered with conventional displays if the prices are set by as few as two or as many as a thousand traders. Providing a suitable tool for market monitoring will not only help market supervisors but also help market engineers analyze and understand the influence of their design.

The crux of the problem is not that there are no advances in technology or new ideas and concepts, but often they are not applied in practice. Even in 2007, scientists encounter financial staff still working with spreadsheet, paper and pencil (Chang et al. 2007). Already in 1996, Matthew Chalmers found that although new products for data analysis were available and traders knew that the new technology could ease their work,

> "their scepticism was kept strong by a stream of over-complex and ill-fitting technology" (Chalmers 1996).

Hence, a new approach has to be designed according to the needs of the user and provide flexible interactions.

Using price and volume information for market monitoring is not enough, because they are too aggregated. They cannot be related to the activity of a single participant. Since fraud is usually caused by a single participant or a group of participants, the *social* context of a participant is relevant when analyzing a transaction. The common visualization of social networks is a graph (e.g. Moody et al. 2005; Heer and Boyd 2005). This can be easily adapted to capture the social characteristcs of the market where participants are represented as nodes and relationships such as HAS SOLD TO are transformed to edges. The problem of common graph visualizations is that after reaching a certain size the graph becomes too large to monitor due to the visual complexity caused by too many elements and the tendency towards visual clutter caused by many crossing edges. The critical size depends on the graph density, size and the applied layout algorithm. Because of the computational complexity, real-time layouting with standard algorithms like the spring embedder is not possible above certain dimensions.

The full transaction and order history that are available on electronic market places contain information about the intentions of the participants. Mining the information that is disclosed to the public is a promising approach to effectively combat fraudulent activities. Since the fraud patterns change whenever fraudsters get aware of a detection mechanism, these patterns have to be reviewed and adapted by humans on a regular basis. The question is how to present the high amount of available data in a meaningful way to the reviewer.

Problem. How can the dynamics of a market be visualized over time while still allowing real time monitoring and not overloading established solutions with semantic complexity?

To avoid ambiguity, in the following, the term *user* shall mean the user of the visualization, whereas *participant* refers to the market participant – also known as *trader*.

The previously mentioned clutter hinders the user's ability to keep track of the active traders, their relationships and the prices and volumes involved. Clutter is mainly comprised of the edges. The contribution of the nodes is negligible. Therefore, several algorithms have been proposed to minimize edge clutter (Purchase et al. 1996; Sugiyama 2002). A typical approach is pre-clustering the data set to make the node placement more intelligent (Archambault et al. 2007). However, a market is a dynamic process and evolves over time. Former neighbors may never trade again. So in order to cluster the current market state in a meaningful way, the trading history has to be considered with decreasing weights for older transactions.

It is difficult to incorporate aggregated information about the market, like price index and overall traded volumes, into the graph. Examples of such complex visualizations can be found in the semiosys software[1] or MatLab[2]. In the following, a combination of a graph and the standard chart displays are proposed to overcome the visual clutter and complexity problem.

3.1 Related Work

Using visualization for market monitoring can be done for two reasons: First, to find opportunities and investment strategies, and second, to reveal fraud and risk. Information visualization in finance has been researched for more than a decade. Most of the publications are dedicated to showing opportunities and depicting the current market situation and development. Fraud detection is only a niche topic. Fraud detection on prediction markets has not been covered yet. Therefore, this section will briefly cover some general publications about market and financial visualizations and then present some closely-related approaches in more detail to compile a requirement list.

[1] S-Explorer `http://www.semiophore.net/v3/en/explorer.html`

[2] MatLab - Biograph `http://www.mathworks.com/applications/compbio/demos.html?file=/products/demos/shipping/bioinfo/biographdemo.html#19`

Various visualizations have been proposed to monitor or compare single stocks against each other. Kiviluoto and Bergius (1997) discuss the use of 2D vs. 3D self organizing maps on financial data. Heatmaps are available and have been used in many trading tools. Treemaps and Information Landscapes have already been mentioned above. Stock Diamonds have been proposed by Blume and Weinhardt (2008) as indicators for recent price development of a single stock.

The application of graph visualization to markets is relatively new. Rostoker (2006) recently published a manuscript on his website regarding interactive presentation of pre-computed stock clusters as a market graph, where the clusters of stocks are represented as nodes and similarities as edges. The graph represents similar price development in different stock groups such as the banking sector, automotive industries, and so on.

In recent research, statistical charts are much better represented. Shmueli and Jank apply box plots and scatter / profile plots to analyze online auctions on eBay (Shmueli and Jank 2005). Besides these more statistical driven approaches, they present scatterplot and list view driven AuctionExplorer in Shmueli et al. (2006). A more statistical approach to fraud detection is presented by Shah et al. (2003). They try to classify transactions and develop rule-sets to distinguish bidding strategies.

Little work has been published so far in the field of market visualization for fraud detection. In the following, a requirement list has been compiled from related works. Some of the articles cover only similar or related topics, such as credit card fraud or visualization of business processes, since the field of market visualization for fraud detection has only recently started.

Ming C. Hao et al. (2006) present an interactive approach for mining business data for fraud. They pre-process the data to determine the driving impact relationship, via correlation and similarity analysis, followed by clustering and classification. The visualization depicts three attributes: source, intermediate and destination. These have been found to be correlated as a cause and effect chain when placed on a circle from left to right. The source attribute is shown on the left half, the intermediate on the vertical diameter and the destination on the right half of the circle. The attributes are grouped in clusters with the lines of the chain connecting the clusters. The main idea is to derive the angle of a node from its importance and thus to place important nodes to the lower area of the diagram. Further emphasis is on the possibility to drill down, filter and mark outliers.

Senator et al. (1995) present the FinCEN AI System (FAIS), that is used at the U.S. Department of the Treasury to mine the data gathered following the Bank Security Act (2000) for money laundering and other types of fraud by using a combination of clustering, link analysis, and visualization. The main application consists of a query interface, which is more comfortable for the user than plain SQL queries. After identifying a cluster or group of interest, the case data set can be exported to the NetMap[3] application

[3] NetMap Homepage http://www.altaanalytics.com/

to be analyzed visually (Goldberg 1998). Besides presenting the application, the authors identify open challenges: better automatic layouts for the visualization of network analysis, identification of "key" nodes, detection of similar sub graphs, (near) real-time interaction, incorporating data from other sources, and finally "temporal link analysis", i.e. introducing simultaneous unlinked activities as new links.

A more stock exchange related approach is presented by Senator (2000). The National Association Of Securities Dealers (NASD) Regulation's Advanced Detection System (ADS) works with a rule engine derived from patterns achieved by mining data from different sources. The mining can be done automatically or by experts. The visualization in ADS is related more to the preconditions and consequences of the rules than visualizing relationships of cases like FAIS does. In a follow-up case study, Senator et al. (2002) describe some of the requirements the system has to meet. Many of the requirements are related to the work of NASD like

> "policy approval and management acceptance of the business use of discovered knowledge".

However, some of the requirements can be seen as more general, like integrated and continuous real-time detection, periodic evaluation and updating of the patterns to keep pace with the fraudsters, and equal treatment of all market participants. Since ADS has no visual analytics part, the requirements are more detection related, but they are still relevant criteria for a market monitoring application.

Finally, looking at business process visualization, Bobrik et al. (2005) compiled a very general list of requirements. They sort the requirements into different groups, namely data integration, user-specific visualization, automatic layout, further requirements, and non-functional requirements. Since some of the requirements have already been mentioned, only new elements concerning the visualization will be listed here: central spot of information, adequate visualization for different user groups with different motivations, highly flexible, generic view concept, different visualization forms, and adaptable graphical parameters.

Closely related to the approach presented in this chapter –though focusing on the mining approach instead of the visualization– is an article recently published by Chau and Faloutsos (2006). They identify potential fraudsters on eBay by analyzing their ratings, transaction history, and their partners. The approach consists of a crawler for the eBay transaction data, a trained C4.5 decision tree, classifying the aggregated transactions into *fraudulent* and *honest*, and a Markov Random Field Model using the network features to distinguishing between three states: *fraudster*, *accomplice*, and *honest* by applying a Belief Propagation Algorithm. The results of this process are visualized as a graph to give the user an overview of the situation. The layout is a standard spring embedder which is transformed manually into a hierarchical layout. The visualization is a report for the end-user, who has queried the trustworthiness of a certain account. Note that the

	Senator et al.	Chau & Faloutsos	Chang et al.
Application			
query interface ‡	☑	☑	☑
incorporate data from other sources†	☑		
"temporal links"‡			
integrated and continuous detection†	☑		
periodic evaluations and update of the pattern†¶	☑	☑	
equal treatment of all market participants†	☑	□	☑
Visualization			
weighted node positions			
drill down$			☑
filtering$			☑
(near) real-time interaction$			☑
automatic layouts$			☑
identification of "key" nodes$¶		☑	☑
similar sub graphs$¶		☑	☑
marking outliers¶		☑	
central spot of information¶		☑	
adequate visualization for different user groups†			
highly flexible, generic view concepts$			□
different visualization forms$			☑
adaptable graphical parameters$			☑

† Senator et al. $Chang et al.
‡ Goldberg et al. ¶Chau and Faloutsos

Table 3.1: Requirement list for a market visualization for fraud detection

approach is based on a set of known fraudsters to train the decision tree in advance.

3.2 Graph Visualization and its Difficulties

3.2.1 Requirements for Visualizations

A list of requirements for a visual market monitoring application has been compiled from the related works. Most of the related work does not directly address financial markets. Some focus on very special aspects; therefore, only the most relevant or most similar approaches have been extracted (see Table 3.1).

These approaches and their limitations are evaluated according to their current published state in the referenced articles.

As the review of related work has shown, most visualizations of markets deal with information aggregation on attributes like price, trading volume, order book, and spread.

However, there is also a social component in the market: the institutionalized exchange of goods. Markets may suffer from collusion or other forms of fraud based on agreements between (sub-) groups of market participants. In order to focus on the social interaction, this work proposes a graph based visualization focusing on the active parties in the market. The questions addressed are:

1. Who is currently active in the market?
2. With whom are they currently trading?
3. Who are their past trading partners?
4. Who is a high volume trader?
5. Who is continuously active and who has only a temporary involvement?

In the following, the visualization will be presented with references back to these initial questions.

3.2.2 Transforming a Market into a Graph

To make the mapping as intuitive as possible, a straight forward solution was taken to construct a graph, which

- maps traders to nodes
- maps sell transactions into directed edges from the seller towards the buyer
- marks the trading volume by the thickness of the edge
- uses an intelligent filtering to focus on the most important events.

This mapping is formally presented in the following definitions.

Definition 3.1. (Vertices)

Every market participant v is mapped to exactly one node (vertex) by a bijective function. The set V of all vertices is defined as follows

$$V = \{v | v \, \mathrm{is\,a\,user\,trading\,on\,the\,platform}\}$$

Definition 3.2. (Share Set)

Further, let S denote the set of all shares traded on the platform.

$$S = \{s | s \, \mathrm{is\,a\,share\,of\,a\,tradable\,product\,or\,contract}\}$$

Definition 3.3. (Transaction Set)

Let $T_{s,t}$ be the set of tuples

$$T_{s,t} = \{(u, v, s, n, p, t) |\, u, v \in V,\ s \in S,\ n, t \in \mathbb{N},\ m \in \mathbb{R}\}$$

representing transactions where u sold n stocks of share s at price m to v at the time t. The time t can be seen as discrete (e.g. milliseconds) and is modeled by a positive integer.

Transactions represent the links (edges) in our graph. Hence, the edge set can now be defined in the following way:

Definition 3.4. (Edge Set)

Let $E_s$be defined as an aggregation of $T_{s,t}$with

$$E_s = \{(u, v, s) |\, (u, v, s, n, m, t) \in T_{s,t},\ n, t \in \mathbb{N},\ m \in \mathbb{R}\}$$

The time t, the amount n, and the price m are aggregated and are therefore not reflected in the edge tuple. They are stored in the attribute set of an edge. The graph can now be defined as

$$G_s = \{V, E_s\}$$

Note: According to these definitions, E_s may contain transactions regarding different shares. Therefore, G_s is a graph representing all activity on the market platform. The user of the application may filter for any desired combination of shares to narrow the investigation.

3.2.3 Modification of the Force Directed Layout

Since there is no inherent physical representation of the trading process on a stock exchange, the general approach is to use the spring embedder layout. The advantage of this approach is that all nodes are treated equally and no hierarchy is superimposed. Early concepts of this layout date back to VLSI design which Eades adapted in 1984 to the field of graph layout (Eades 1984). A good introduction can be found in Brandes (2001). Following Brandes, the algorithm is based on two force functions and an update loop:

Definition 3.5. (Force functions):

Let p_x denote the vector pointing to the position of node x on the drawing space. Then the force functions can be defined as follows:

Algorithm 3.1 Basic spring embedder algorithm taken from Brandes (2001)

```
input: connected undirected graph $G(V, E)$,
       an initial placement vector p
output: placement vector p with low internal stress
for $t \leftarrow 1$ to $iterations$ do
  # calculate forces on nodes
  for $v \in V$ do
    $F_v(t) \leftarrow \sum_{u:(u,v)\notin E} f_{rep}(p_u, p_v) + \sum_{u:(u,v)\in E} f_{spring}(p_u, p_v)$
  # update node positions
  for $v \in V$ do
    $p_v \leftarrow p_v + \delta * F_v(t)$
```

$$f_{rep}(p_u, p_v) = \frac{c_g}{\|p_v - p_u\|^2} * \overrightarrow{p_u p_v} \tag{3.1}$$

$$f_{spring}(p_u, p_v) = C_\sigma * \log \frac{\|p_u - p_v\|}{l} * \overrightarrow{p_v p_u} \tag{3.2}$$

The function (3.1) represents repelling forces between unrelated nodes, whereas the function (3.2) represents spring forces between nodes with a common edge. The repelling function (3.1) returns the vector of the direction in which node u pushes node v away ($\overrightarrow{p_u p_v}$), weighted by a gravitational constant c_g related to the current distance squared. The spring forces (3.2) between two nodes with a common edge depend on the current distance divided by the desired spring length l. The logarithm of this term guarantees that the contracting force is positive if the current length is greater than l, negative if lower than l, and zero if the length equals l. Finally, the force is weighted by a spring coefficient c_σ and drags node u towards node v by the vector $\overrightarrow{p_v p_u}$.

These two force functions are combined in an update-algorithm which iteratively calculates all influencing forces on each node (see Algorithm 3.1). The input is a connected undirected graph and a vector with an initial placement of all nodes (i.e. randomly chosen). The first inner for-loop calculates the forces on all nodes. The second inner for-loop updates the node positions. These two steps cannot be done in a single loop. Otherwise nodes that are processed later in the for-loop would already take updated node positions into account.

Since the first inner for-loop iterates over all nodes in which the force-functions again iterate over all nodes, the run-time complexity is of $O(n^2)$ for each iteration. Several solutions have been proposed to speed up the algorithm. Fruchtermann and Reingold (1991) propose squared force functions to speed up the initial positioning of the nodes . Barnes and Hut (1986) present a hierarchical approach applying a divide and conquer pre-

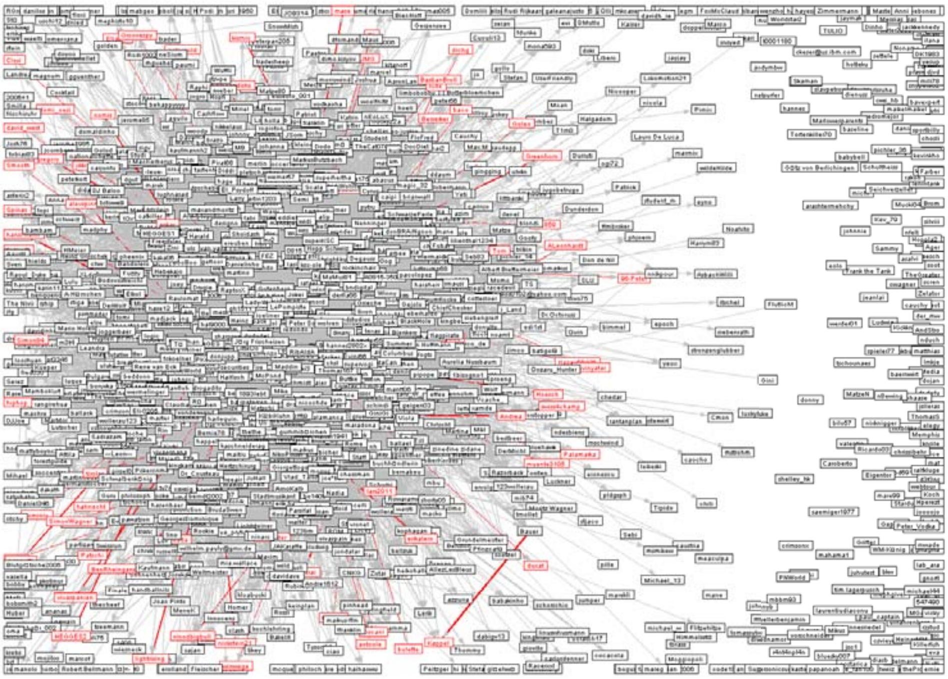

Figure 3.1: Example clutter when visualizing a market with ca. 1500 nodes using the standard spring embedder

calculation and thus reduce the overall run-time complexity from $O(n^3)$ to $O(n^2 \log n)$ with an acceptable error degree. Their approach is implemented in the prefuse[4] framework on which the presented prototype is based.

The algorithm presented above is defined for an undirected graph. Although the previously defined edge set of a market contains directed edges (a has sold to b), the layout algorithm works fine. The force function iterates over all nodes and thus takes all edges into consideration.

Applying this algorithm to market data will sooner or later lead to visual clutter due to the arbitrary relations between market participants. Figure 3.1 depicts a market with 1,543 traders and 34,748 transactions. Trades happening on the current day as well as new traders joining the market on the current day are highlighted in red color. But because of the visual clutter and compactness it is hard to grasp what is currently going on.

To improve the visualization, the algorithm needs to be modified. Since market graphs vary in their density and evolution, the spring embedder has to adapt to the current market situation. The goal is to highlight the current developments in the context of the previous trading situation. In addition, the visualization should allow the user to follow the market activity without the need to scroll, zoom, or pan. Innovative visualizations from the field of network monitoring, which fulfill these requirements, are proposed by Livnat et al. (2005a,b). They describe a circular visualization, in which the current events

[4] prefuse visualization toolkit `http://prefuse.org`

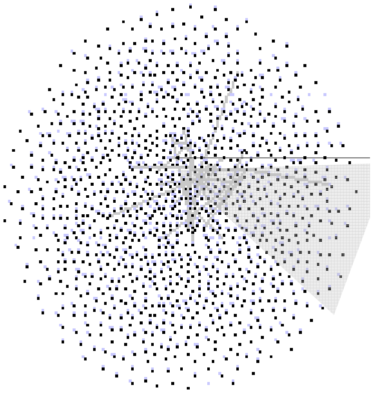

Figure 3.2: Market graph mash-up; the gray triangle indicates the history-area of the central nodes. Inactive one-time traders get pushed out by new one-time traders within the gray area

appear in the center, while past events move slowly outward, creating circles like annual rings of trees. The circles are divided into sections related to a certain event type which determines the drifting direction of each occurring event.

Creating a similar metaphor for markets might help the market operator to understand the duration of social relationships between participants. Figure 3.2 shows a map overview of a possible visualization. The graph is embedded in the center of the visualization. The edges are shown as gray lines. The dots symbolize formerly active accounts. They are shown without edges since they did not trade within the time slice of the visualization. Active nodes are dragged towards the center of the visualization driving inactive nodes to the outer area. Thus, the most outer belt contains the accounts that have been inactive for the longest time. An activity history is inherent in the structure and is indicated by the gray triangle and the time line. Note, since the repelling forces diminish with the distance of the nodes, all nodes without edges relax to a default distance between each other.

To structure a spring embedder visualization as described above, the animation and update functions have to be fine tuned to the typical relationships occurring in markets. This modification will be described in detail in the following subsections. For a better understanding where the modifications apply, a brief introduction of the underlying model will be given first.

Figure 3.3 shows a sequence diagram of the application. The market platform operates directly on the database. The visualization prototype constructs a model of the market from the database. Some filters are applied, which leads to a second reduced model. This information is provided to the layouter, who assigns coordinates to each visual object and

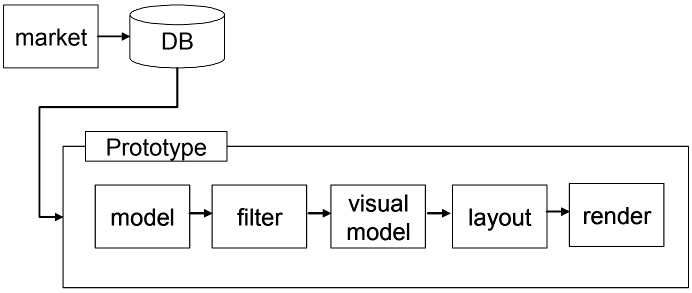

Figure 3.3: Sequence diagram of the different steps in the prefuse visualization chain

finally the renderer decides the appearance of each object according to its class and attributes.

3.2.4 Filtering the Edge Set

The first challenge is the growing number of edges and nodes over time. A moderately sparse graph with more than 50 participants and 200 transactions already contains too much information to be analyzed in real-time. This is a problem of information overload. To reduce the amount of information presented, a sliding window can be applied to the transaction set reducing the number of edges drawn simultaneously. Only transactions concluded in the last n time periods (e.g. seconds) are displayed as edges of the graph. Edges based on older transactions are removed from the visualization. In order to be able to filter the graph in the desired way, a timestamp attribute is introduced into the edge data structure (see Appendix A). However, nodes - once they are created - remain in the layout throughout the visualization. Though removing the nodes from the visualization would further speed up the layout process, they still indicate previous business relations with neighboring nodes.

For example, imagine a case where one node is surrounded by fraudsters and all edges are already filtered. A market operator would pay more attention to this account than if all the surrounding nodes are already removed. Therefore, it was decided not to filter the nodes in the visualization.

Definition 3.6. (Filtered Edge Set)

Let $E_{s,\Delta t}$ be defined as

$$E_{s,\Delta t} = \{(u,v,s,t) | \, (u,v,s,n,m,t) \in T_{s,t}, \, n,t \in \mathbb{N}, \, p \in \mathbb{R} \land t_{now} - t \leq \Delta t\}$$

The edges of the set are based on transactions in the share list s concluded *shortly before* or *right on* the current time step t_{now}. So the resulting graph $G_{s,\Delta t} = \{V, E_{s,\Delta t}\}$ allows the user to filter by time and share.

Filtering the edge set by time is just one possibility. It answers the first and second of the initial questions: who is currently active in the market and with whom they are trading?

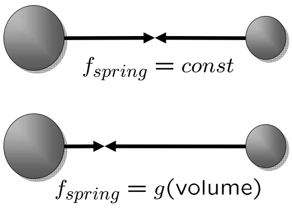

Figure 3.4: When introducing mass $\sim$ trading volume (depicted by the node size), the new spring force -dragging the nodes towards each other- moves the smaller node faster, and thus closer to the high volume trader

Further filtering options result from the other attributes stored in the transaction set. One could filter for price, volume, share, or user. The prototype presented in chapter 4 only supports filtering by time and share.

3.2.5 Introducing Mass

Since filtering only reduces the complexity, the representation so far only provides who traded with whom and when. Interesting aspects such as who is a heavy and/or active trader are hardly recognizable from the animation. The layout algorithm forces the nodes to jump together whenever a visual edge is inserted in between them and drift away after the edge disappears. To keep the attention of the user on a central spot and maintain his mental map, a semantic has to be introduced into the node movement behavior. To distinguish heavy traders from low volume traders, a mass (also interpretable as a drag coefficient) related to the overall trading volume of the node is introduced. An example is given in Figure 3.4. Instead of equal spring forces, the high volume trader, who is depicted with the bigger circle, will move slower and therefore less towards the low volume trader than the low volume trader will move towards the high volume trader.

The overall trading volume of each market participant is the sum of all his buy and sell transactions:

$$\mathrm{vol}(v \in V) = \underbrace{\sum_{e:(u,v,s,n,m,t) \in T_s} n_e * m_e}_{buy\,volume} + \underbrace{\sum_{e:(v,u,s,n,m,t) \in T_s} n_e * m_e}_{sell\,volume} \tag{3.3}$$

Empirically, a exponential distribution of the volume could be observed in data sets of the stock exchange as well as in prediction markets (see Figure 3.5). A non-linear statistical regression underlines this relationship (see Table 3.2). Therefore, a logarithmic transformation of the total trading volume is used as a drag coefficient. The update loop

$c\exp^{\lambda x}$- *Regression*	R^2	c	λ
Bundesliga 1 06/07	0.894	2090909	-0.81
Bundesliga 2 06/07	0.810	2558904	-0,123
Bundesliga 05/06	0.931	837754,7	-0.43
STOCCER WM	0.774	2429544	-0.006
Canada	0.869	51156810	-0.94

Table 3.2: Exponential regression analysis, underlining the hypothesis of exponential distributed trading volumes on financial market places by high R^2 values

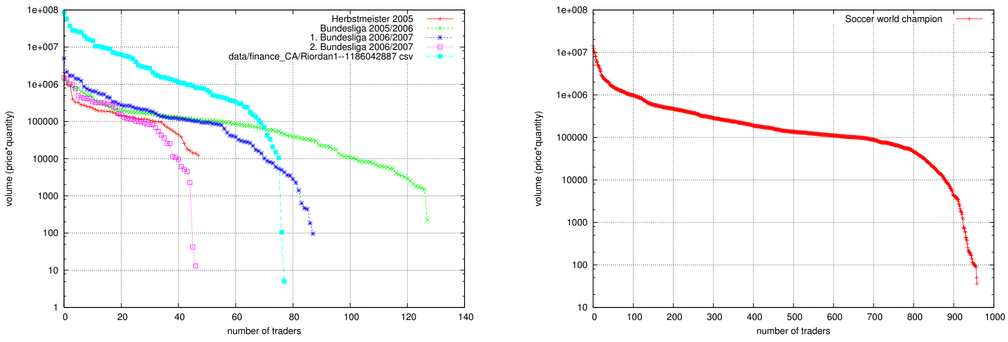

Figure 3.5: Total trading volume per participant in different markets

of algorithm 3.1 has to be modified in the following way:

$$p_v \leftarrow p_v + \frac{F_v(t)}{\delta(v)}$$

to incorporate the drag coefficient function $\delta(v)$, which is defined as

$$\delta(v) = c_\delta * \log_{10}(\mathrm{vol}(v))$$

where $c_\delta$is a constant for calibration.

Less active traders, who have a low total trading volume, now slowly disperse from the center of the visualization. When they transact they are 'snapped' closer to the center of the visualization depending on the prior trading volume of their partners. The more active a trader is, the slower he will move. For example, a heavy trader will stay almost stationary in the middle of the visualization, this being a good orientation for the observer of the visualization. This helps the user answer the fourth initial question of who is a high volume trader.

New nodes entering the market are placed close to the center of the visualization; the exact location is determined by its edges and trading partners. Because of the repelling forces older nodes are pushed to the outskirts. This causes the graph to grow in concentric circles, where the inner circles represent the most recent active traders (see Fig. 3.2).

One time traders move out of the center faster than traders coming back every now and then. So the visualization is continuously updated corresponding to the trading activity.

3.2.6 Adjusting the Repelling Forces

The third initial question regards the trading history of a node. Therefore, the next modification is related to the problem caused by repelling forces leading to a uniform distribution of the edgeless nodes in the outer part of the visualization. These belts of inactive traders provide little information because their placement is only influenced by the repelling forces, which are equal towards all other market participants.

Since the visualization should answer who was formerly trading with whom, the placement should be an intuitive representation of this relation. The full answer to this question is a $|V|$-dimensional vector, which requires a $|V|$-dimensional space; however, the visualization is only in a 2-dimensional space. In order to properly depict a market with more than two participants, the requirement have to be relaxed. Instead of answering this question for all nodes, the visualization only shows the relation between visual neighbors. If two nodes have the default distance, they have never traded with each other. If they are closer together, it is because they have traded together before.

To reach a meaningful clustering, the repelling force (3.1) has to be modified. The idea is to reduce the repelling force between two nodes in case of former trading activity. At first sight, this looks like having to traverse all other nodes in each repelling force calculation again. So, the second objective is to keep the necessary extra computation as low as possible. This can be achieved via a hash look-up in the full, unfiltered edge set (E_s). The function is defined as follows:

$$\rho_{cluster}(p_u, p_v) = \begin{cases} 1 & \text{if } (u, v, \dots) \in E_s \\ \epsilon & \text{if } (u, v, \dots) \notin E_s \end{cases} \tag{3.4}$$

with $\epsilon \in (0, 1)$. Empirically, good results have been achieved setting $\epsilon = \frac{1}{8}$. The cluster function $\rho_{cluster}$ is now introduced to reduce the repelling forces:

$$F_v(t) \leftarrow \sum_{u:(u,v) \notin E_{s,\Delta t}} \rho_{cluster}(p_u, p_v) * f_{rep}(p_u, p_v) + \sum_{u:(u,v) \in E_{s,\Delta t}} f_{spring}(p_u, p_v)$$

Note that $\rho_{cluster}$ depends on the full edge set, while operating on the filtered edge set. Thus the visualization knows about the edges, although they are not painted anymore.

3.2.7 Adjusting the Edge Width

The next modification improves the rendering. For the market operator, it is highly desirable to distinguish high and low volume trading relationships. A very intuitive reflection

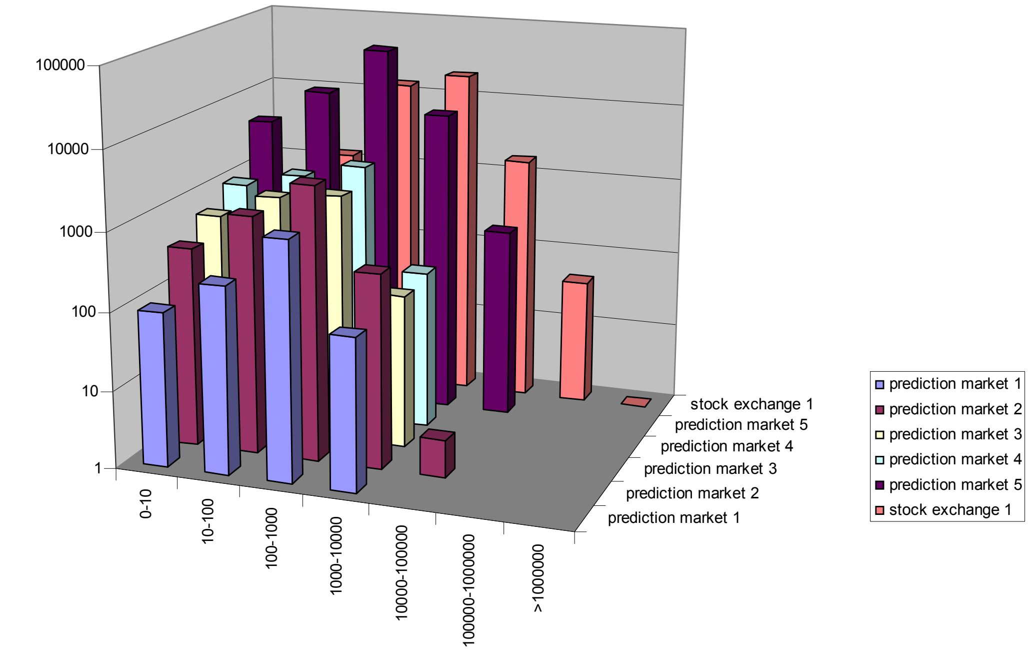

	0-10	10-100	100-1000	1000-10000	10000-100000	100000-1000000	>1000000
■ prediction market 1	96	253	1117	88			
■ prediction market 2	367	1134	3203	310	3		
■ prediction market 3	606	1236	1502	95			
■ prediction market 4	1013	1540	2334	110			
■ prediction market 5	4778	13448	53375	8519	282		
■ stock exchange 1	96	1274	13578	19945	1629	43	1

Figure 3.6: Transaction volume histogram

of the trading volume can be represented by drawing the edge width according to the volume (price * quantity) traded over this edge.

The transaction volume distribution analysis is depicted in Figure 3.6 for data sets of five sports prediction markets and one financial stock market. The histogram classes show the number of shares traded in each transaction. For example the first class represents the number of transactions where 0 to 10 shares have been traded. Note that both the histogram classes and the y-axis have a logarithmic scale. Common within all prediction markets is a peak in the 100-1000 shares class. Only the financial stock exchange data set has a peak in the 1000-10000 class. The order volumes reach up to a million and more in the financial market.

To be able to distinguish the different empirically observed classes in the visualization, the edge width function uses the logarithm base 10 of the volume. The flexibility requirement implies that the visualization has to adapt to each market automatically. So in low volume markets, the full spectrum of edge widths should still be applied. Furthermore, an edge width above a certain limit is perceived to be rather unaesthetic. To solve these two aspects, the volume is normalized with the maximum volume traded over an edge so far. Since the normalized scale range of the standard spring embedder algorithm is $[0, 1]$, the result has to be scaled again to $[1, 7]$. These numbers are not arbitrary but related to the seven empirically observed classes and the observation that pixels are still esthetically acceptable. If a number larger than seven is selected, the nodes are only a little larger than the edges and thus hard to distinguish.

$$width(e \in E_s) = 1 + 7 * \frac{\log_{10}(1 + e.volume)}{\log_{10}(1 + \max i.volume)} \tag{3.5}$$

The edge width function (3.5) adds 1 to the volume to prevent negative logarithm values and thus a negative width. Since the fraction only scales to the range of $[0, 6.1)$ and an edge width of 0 is not desired, the overall value is incremented by one. The resulting function is depicted in Figure 3.7. The x-axis represents the maximum volume observed in the market so far, while the z-axis is the current edge e. The y-axis is the value of the function $width(e)$. It can be seen that the function is linear so far in the case of the observed maximum of $1.000.000$ on the x-axis. When going towards 1 on the x-axis, the function becomes more logarithmical because the current calculated value of e is simultaneously the current observed maximum and therefore stretches the linear starting scale.

This behavior may not always be desired. If the maximum volume of the last evaluated edge is significantly higher than all volumes observed before, the visualization will be misleading. All other edges, which were evaluated with the old maximum, have been drawn too thick in comparison to the last edge. This can be handled in two ways:

1. determining the maximum in advance, for example, on creation or when updating the graph
2. ignoring it, since the visualization in this case is continuously updated like a video, typically 25 fps, meaning the error will only be visible for 40ms.

In the prototype presented in the next chapter, the second option has been chosen. The reason for this choice is that, in the model-view-controller concept, the model and its knowledge is separated from the view classes. In the prefuse toolkit, the edge-renderer class does not have access to the full graph and its information but only to the edge it is currently rendering.

An example screenshot of the visualization can be seen in Figure 3.8. The central players are easy to recognize and their trading relationships are clearly visible. The graph is not necessarily connected. For example, node 125 is selling to 83 in the current time slot Δt but neither of them is trading with anyone else. In the middle of the lower part, the two nodes 90 and 107 are closer together than the default distance. This indicates that they had been trading with each other previously. Another example of this are nodes 74 and 43 on the right. The inactive nodes form belts around the center of the graph. The most inner unconnected belt is the group of traders who just stopped trading.

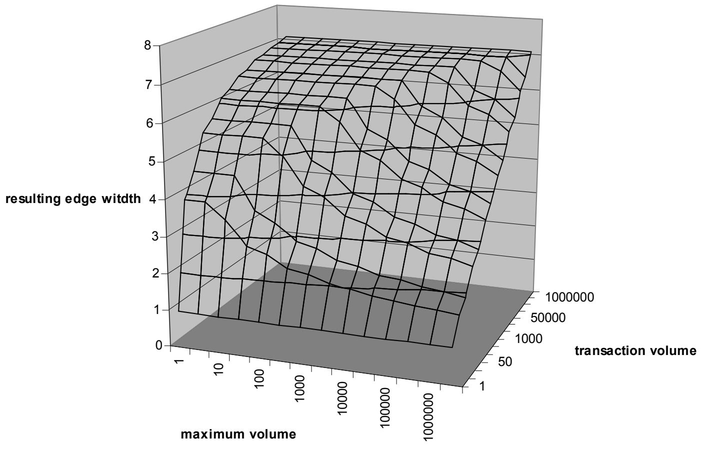

Figure 3.7: Edge width function (y) with varying transaction and maximum volumina

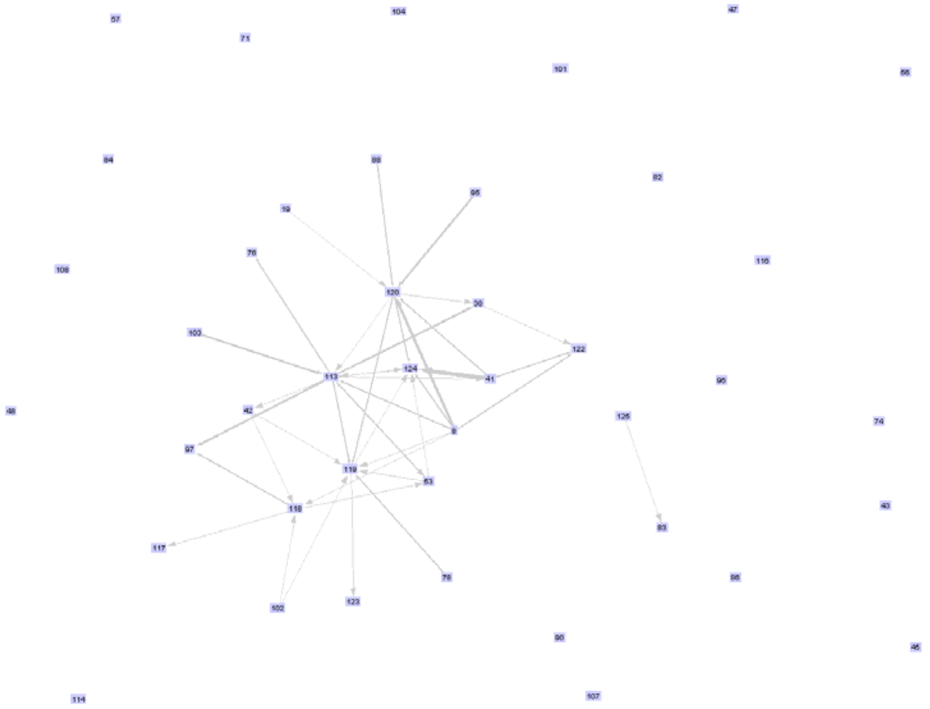

Figure 3.8: Example for volume related edges

3.2.8 Combining Visualizations for the Application

The graph itself is a very intuitive visualization of the social relationship among market participants. But besides these patterns, investigators need an overview of the current development of the market. This is necessary to relate the observed market events with social entities and patterns. Market events are such things as price jumps or uncommonly high or low trading volumes of a certain share.

To keep the visualization as simple as possible, these financial aspects regarding the whole market are not incorporated into the graph visualization. However, the typical charts, like bar charts for trading volume and scatterplots for the price development, are presented at the side of the graph visualization (see Figure 4.14 on page 84). The scatterplot for the prices displays the price index by default. If a peak is observed, the user can drill down to the scatterplot of a single share via the drop down list of the combobox. The volumes are presented as stacked bar charts. Each share is represented in a different color. On the upper right, a map view of the current graph is displayed.

3.3 Evaluation

There is no standard procedure to evaluate a visualization as it may serve different needs. Three main aspects of visualization evaluation will be presented in the following:

1. complexity of the underlying algorithms
2. human-computer-interface perspective
3. user requirement fulfilment

In the following sub sections, the above aspects will be explained and applied to the proposed visualization.

3.3.1 Algorithmic Complexity

The standard way to measure algorithmical complexity is by giving upper or lower asymptotic bounds in the big-o notation. The o notation was introduced by the mathematician Paul Bachman (1894). Public attention to this notation was drawn by the German Edmund Landau (1909), which is why the notation is also known as Landau symbols. Nowadays, the usage is widespread in standard computer science literature (e.g. Knuth (1997); Sipser (1997)). In graph theory, where the graph is typically denoted in $G = \{V, E\}$, the problem dimension is separated into the number of nodes $|V|$ and the number of edges $|E|$.

The upper boundary of a standard spring layout as presented in Alg. 3.1 on page 54 is in $O(|V|^3)$. One approach to reduce the runtime complexity is pre-clustering the nodes

Algorithm 3.2 Final version of the spring embedder including all modifications

```
input: connected undirected graph $G(V_s, E_{s,\Delta t})$,
       placement vector $p$
       $E_s$-hashtable for the cluster function $\rho_{cluster}$
output: placement vector $p$ with low internal stress
for $t \leftarrow 1$ to $iterations$ do
  for $v \in V$ do
    $F_v(t) \leftarrow \sum_{u:(u,v)\notin E_{s,\Delta t}} \rho_{cluster}(p_u, p_v) * f_{rep}(p_u, p_v) + \sum_{u:(u,v)\in E_{s,\Delta t}} f_{spring}(p_u, p_v)$
  for $v \in V$ do
    $p_v \leftarrow p_v + \delta(v) * F_v(t)$
```

and grouping them in a quad-tree (Barnes and Hut 1986). After a certain distance, the virtual nodes representing a cluster of nodes are only used for calculation. This reduces the runtime complexity to $O(|V|^2 \log |V|)$. Due to the notational complexity, only the default spring embedder with the previously described modifications (see Alg. 3.2) will now be considered. Whether the assumptions also hold true for the Barnes-Hut approach will be discussed for each modification.

The filtering reduces the number of times f_{spring} has to be calculated in the first for-loop. But instead of calculating the spring force to the neighbor, now f_{rep} has to be evaluated for this node. Since the difference in complexity between f_{spring} and f_{rep} is negligible, the filtering is neutral in terms of runtime complexity.

The second modification in the first for-loop is the introduction of $\rho_{cluster}$. The cluster function consists of a look-up whether an edge exists between u and v. This look-up can be done with a hashtable in $O(1)$ and therefore can be considered neutral in terms of runtime complexity.

Barnes-Hut reduce the number of nodes which f_{rep} is summing up by introducing virtual nodes. Since these virtual nodes are not reflected in the hashtable of $E_s$for the cluster function, $\rho_{cluster}$ will return the default distance coefficient 1 (c.f. equation 3.4). To estimate the error introduced into the clustering, the virtual node creation and selection has to be explained in more detail. The quad-tree is constructed via partitioning the space into rectangles and splitting them repeatedly until there is only one node left per rectangle (see Figure 3.9). In the depicted example, it can be seen that only nodes for non-empty rectangles are added to the tree. The virtual nodes are initialized in the barycenter of the child nodes. The decision whether to use the virtual node or traverse the tree to the child nodes can be expressed by the following condition:

$$\frac{\textit{size of rectangle}}{\text{distance of barycenter to current node}} < \beta$$

This term can be read as "if the rectangle is far enough, we can safely use the virtual node". For our cluster function $\rho_{cluster}$, this is irrelevant since the only goal is to reflect

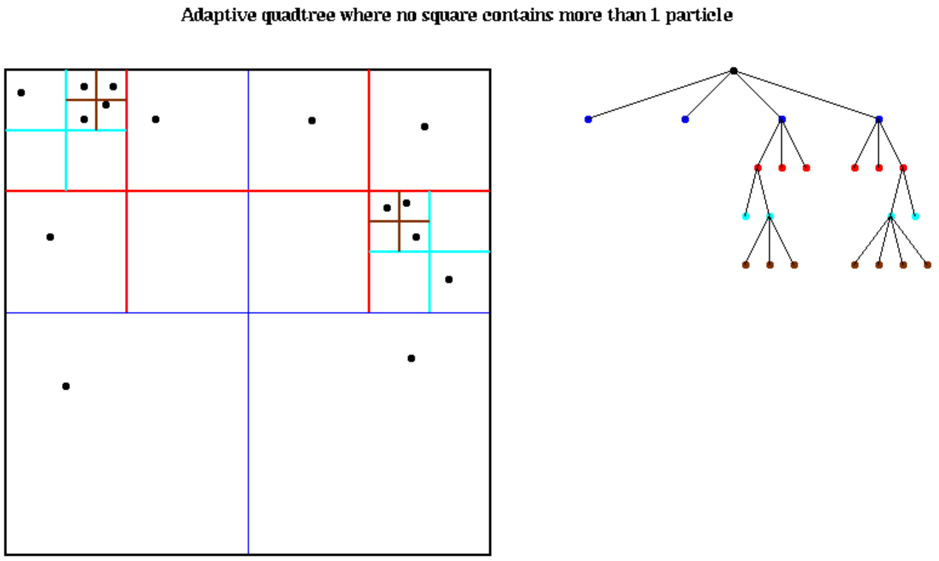
Adaptive quadtree where no square contains more than 1 particle

Figure 3.9: Quad-tree transformation in the Barnes-Hut reduction (Dehmel 1996)

the relationships in the nearest neighborhood. As the condition only selects virtual nodes for larger distances, there is no error introduced by this modification.

In the second for-loop, inertia is introduced via the function $\delta(v)$. This is only a linear transformation of a node attribute. Since the node attributes are calculated upon creation or update of the graph and the function $\delta(v)$ is only a linear transformation, $\delta(v)$ is element of $O(1)$. Clearly, this also holds in case of the hierarchical approach.

The last modification concerns the edge width. This is only a rendering aspect; and furthermore, only a linear transformation of the volume attribute ($width(e \in E_s) \in O(1)$). Thus, it is not relevant for the layout algorithm at all.

Summarizing the influence of the modifications on the runtime complexity, it has been shown that there is no significant change in the complexity. The modifications can be applied to the basic algorithm as well as the hierarchical approach without significant performance loss. However, the memory complexity grows from $|V| + |E_s|$ to $|V| + |E_s| + |E_{s,\Delta t}|$ due to the look-up table for $\rho_{cluster}$ and the filtered edge set for the visualization. Though the size of $E_{s,\Delta t}$ depends on Δt, it can be safely stated that $E_{s,\Delta t} \subseteq E_s$ usually with $|E_{s,\Delta t}| \ll |E_s|$.

3.3.2 User Requirement Fulfilment

There are several ways of measuring requirement fulfilment. In the context of ergonomics, it can be done by using a participating observation, requirement analysis, checklists, and the like. In the field of ergonomics, many descriptions have been applied to describe more or less the same concept, referred to here as human-computer-interface (HCI). Software ergonomics, human-computer-interaction, GUI design, and usability engineering are often

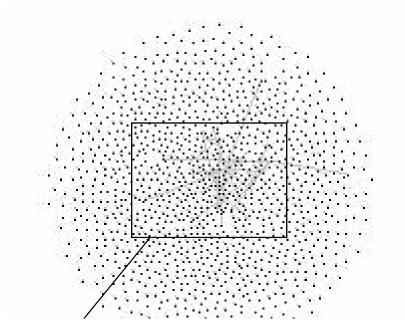

Figure 3.10: The clipping or visual area is markedby the black rectangle as compared to the size of the graph

discussed, though they describe different concepts. Software ergonomics / HCI principles are summarized in the ISO norm 9241, parts 10-17 and 110 and multimedia systems are summarized in DIN EN ISO 14915. Most of these concepts refer to software systems. HCI principles and the requirement list that were gathered in the beginning of this chapter will be discussed in detail after the presentation of the prototype (c.f. Section 4) in chapter 5 on page 97. In this chapter, the visualization is evaluated according to the goal of presenting a compact visualization for monitoring purposes.

The spring force visualization is not necessarily deterministic. This holds especially true when the user can interact with the visualization. A good visualization should point the user towards the interesting aspects of the underlying data. Especially in the case of monitoring a continuous task, the visualization has to adapt to the market situation and should be readable without zooming or scrolling. Hence, setting the zoom level far enough out of the visualization would solve the problem. If the user has to zoom and scroll continuously, it distracts or even hinders him from judging the situation and taking appropriate action.

The question is how to evaluate a visualization while being mindful of these requirements? There is no theoretical approach to predict the space consumption a spring embedder layout will need. There is only a worst case upper boundary. In the worst case, the nodes will form a long chain, each node connecting with exactly two neighbors, except the first and the last node. This chain would relax into a long diagonal with the length of $(|V|-1)*l*c_\sigma$. But besides this worst case scenario, no further predictions about the space consumption can be made. Thus, only empirical evidence can be given. A straight-forward approach to gathering such empirical data is to count the edges drawn inside the visual clipping area and then count the number of edges drawn outside. Each edge drawn outside of the visual clipping area would require the user to scroll or zoom to see this transaction. This metric will be defined as the percentage of edges drawn outside the clipping area. Let G be a set of graphs, then:

$$metric(G) = \frac{\#edges\,drawn\,outside}{\#edges\,drawn\,inside + \#edges\,drawn\,outside},$$

An open question is which data set to use for such a metric. Since there is no default market and the only explicit requirement for the visualization is the flexibility to adapt to different markets, it is hard to test its robustness. There are two possibilities: Generate different graphs to simulate different markets or take the graphs of existing market data sets.

For the first approach, it would be necessary to cover all possible extreme settings to challenge the visualization. But it would still be rather doubtful whether the employed settings sufficiently reflect reality. Theoretically, there can be three reasons for edges to be drawn outside the clipping area:

1. Unconnected Graph:
   There is a long period with no trading activity in a share and the nodes corresponding to the current spread limit offers are floating outside of the clipping area. If another inactive trader returns to the market and trades against the spread, bying or selling less than or equal to the number of shares offered by the other inactive trader, the edge will be drawn in between them. This edge will be invisible to the user only if the position vectors of both nodes are pointing to the same side of the visualization. Else the edge will cross the clipping area. This case will happen in only 1/4th of these cases.

2. Connected Graph:
   If many members start trading within Δt, it is possible that the filtered graph can grow beyond the clipping area.

3. Graphs, related to the worst case scenario:
   Any graph that contains a longer chain as described above in the worst case scenario -whether connected or unconnected to the rest of the graph- could display only the clipped part, if the chain connects nodes inside and outside the clipping area. This can only happen if *a* sells to *b*, who buys or sells only to *c*, who buys of sells only to *d* and so on, where -without loss of generality- *c* and *d* are outside the clipping area. Though it is not impossible, it seldom occurred in the observed markets (for exact figures see below). This probably relates to the fact that 80% of the trading is done by 20% of the market participants. Less active participants typically trade against the active traders. To form such a chain the inactive traders have to coordinate a process sequence of buying exactly the same size offered by the other player.

The number of edges drawn not only depends on the graph but also on the filter setting Δt. The sliding window setting was altered between one hour and one day. The market ran over 56 days. As the time period grows longer, more edges will be displayed. This makes the layout more complex. The update increment was set to the same value as the sliding window. This means that all edges from the previous update are filtered out for the next update and only the new edges are displayed. This is more challenging for the layout algorithm since improved node placement, in order to reduce edge crossings in the previous step, may be contrary to the optimal placement for the current update. Since force directed layouts continuously improve the node placement in each cycle and can be run to infinity, the relax time was set to six seconds. Thus the layout algorithm had six seconds after each update to improve crossings and node placement of the graph before the time was incremented by Δt and the next update started. After each run the initial node setting was randomly altered.

Overall the visualization ran for ten hours on a 1600x1200 resolution (no zoom, font size at 10pt). On average, only 2.1 % of the edges were drawn completely outside the clipping area with a standard deviation of 1.6. The maximum was 4.9 % in a run with $\Delta t = 1\, day$. In a market with 1500 transactions on average per day, stepping day-by-day through the data set still results in to many relations to get a clear picture. If these less realistic outliers are excluded, the average decreases to 1.4 %. Thus 98.6 % of all activity can be monitored without any scrolling, zooming or panning.

3.4 Summary

In this chapter, a visualization for the social aspects of market activities has been presented. The visualization is a modification of an existing spring layout algorithm. It allows the user to focus only on the important nodes and activities. The runtime analysis shows that the modifications do not significantly increase the complexity (still in $O(n^2 \log n)$) and the empirical test indicates that the visualization accommodates common market activity very well within the dimensions of a common monitor. The interactive capabilities -important for the investigation- are presented and discussed in the next chapter.

4 Software Prototype

The prototype was implemented with the goals of portability, flexibility and intuitive user interface design. The primary user groups addressed are market operators and investigators. Applications in the field of fraud detection span from simple query interfaces to sophisticated link analysis tools. A general distinction can be made between scientific prototypes and reports about real systems. As Bolton and Hand (2002) state in their review, there are no detailed published methods of real systems or any publicly available data sets used to train and evaluate real fraud detection systems. This shall prevent fraudsters from learning too much about the current state of the art. The same holds true for software available for fraud detection. Most products give only a very superficial description of its capabilities and technology.

The *FinCEN AI System* (FAIS) and the *advanced detection system* (ADS) from the *NASD* have already been presented in the previous chapter about the visualization design. Though they are well established in the industry, some scientific publications can still be found about them. Further examples of commercial fraud detection software can be categorized into different application areas such as finance, telecommunications and networks intrusion, credit card, and money laundering. A detailed overview of these examples is given in Appendix E on page 141. In the field of financial markets, prominent examples are *DeltaMaster* (Bissantz & Company GmbH 2008), *Coral8* (2008), *SMARTS* (Smarts Group International Pty Ltd 2008), *Apama* (Progress Software Limited 2008), and mlds (Innovations GmbH 2008). However, only a little information is publicly available about most of these tools. *SMARTS* and *Apama* are event based systems which allow the user to define rules that will be activated on certain events. A more visual approach is made possible by *mlds* which integrates *VisualRules*. This tool allows the user to define and manipulate the rule set in a graphical rule editor. The underlying decision tree of the rule is graphically depicted and can be reordered and parametrized. *Coral8* and *DeltaMaster* are data-mining tools that allow the user to apply certain tests or algorithms. Less specific are pure data mining and rule engine applications like *WEKA* (Witten and Frank 2005) and *GritBot* (RuleQuest Research Pty Ltd 2008). Although the pure data mining applications are powerful and flexible, they require a sound knowledge about data mining. The last empirical evaluation of commercial fraud detection systems was done by Abbott et al. (1998).

A very comparable system, though not based on a graph visualization, but also using a multi-display approach has been recently published at the IEEE Visual Analytics Science and Technology Symposium 2007 by Chang et al. (2007). They describe a system developed in collaboration with *Bank of America* for transaction analysis. It combines a keyword network view, a heatmap, a search-by-example tool, and a newly developed view called *Strings and Beads*. The goal was to facilitate the search for suspicious accounts and transactions especially in the context of terrorism. Thus, interactive filtering over time, keywords, and accounts are of crucial importance for the investigators and are specifically addressed by the interface.

Since portability is a desired feature in the heterogeneous university environment, the prototype underlying this work has been developed in *Java* and *ANSI SQL 2.0* (ISO 9241 1996; Cannan and Otten 1993). The application has been tested on *Windows* and *Linux* in combination with local and remote database servers (*PostgreSQL* 8.0, 8.1, and 8.2).

The chapter is ordered as follows: Section 4.1 presents the use cases of the application and how these are addressed in the prototype, Section 4.2 goes into detail about how the detection algorithms are integrated into the prototype, Section 4.3 explains the user interactions with the visualization, and Section 4.4 presents further investigation tools and helpers. The chapter concludes with a short implementation and design overview.

4.1 Use Cases

The requirements of the prototype can be derived by a use case analysis. Figure 4.1 depicts the central use cases of the prototype. The analysis tasks can be split up into diverse use cases. Each use case will be described in detail in the following section.

4.1.1 Monitoring

The market operator needs continuous feedback about the current development of the platform: e.g., user statistics, price development, transaction volumes, and the interaction of the different market participants. Looking at the minimal market model (Rolli et al. 2004), the enumerated entities reflect the basic entities of markets. Besides tabular lists, graphical elements like bar- and pie-charts are useful to give an oberview. Views should update in realtime. This following definition captures all these aspects on an abstract level:

Definition. Monitoring: Continuous surveillance/observation of a process..

Monitoring is supported in the prototype by the graph visualization described in Chapter 3. The integration of the graph visualization into the prototype and further involved views are explained in more detail in Section 4.2.3 on page 83.

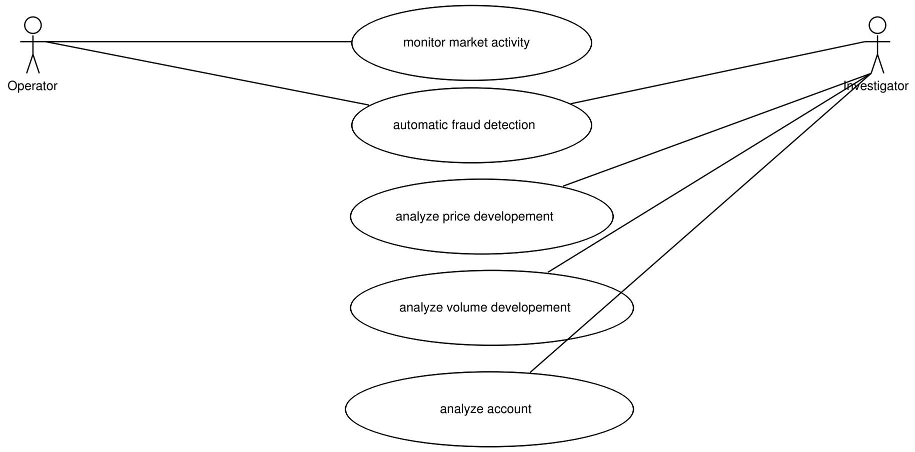
Operator
Investigator
monitor market activity
automatic fraud detection
analyze price developement
analyze volume developement
analyze account

Figure 4.1: Use case diagram

4.1.2 Automatic Fraud Detection

Since the purpose of the prototype is to support the user with semi-automatic fraud detection, the visualization and the designated detection algorithms should incorporate necessary functionality for both aspects.

The visualizations have to allow marking, cross linking and drilling down functionalities. Marking is necessary to follow certain accounts or to mark a special market situation. When deciding whether a situation is relevant or not, it often depends on the context. This can be done via cross linking. The context should be available to the user within only a few clicks. An example of context is looking up a participant's portfolio, when investigating a suspicious transaction, to understand the motivation they had to sell or buy in that specific moment. Drilling down allows the user to disaggregate from the highest level of aggregation down to the most detailed view of information (e.g. offers, transactions, users).

Making the parameters of the detection algorithms available allows the user to interactively adjust the sensitivity of the detection algorithm to the current market situation. At the same time, the number of parameters has to be minimized in order to reduce the complexity for the user.

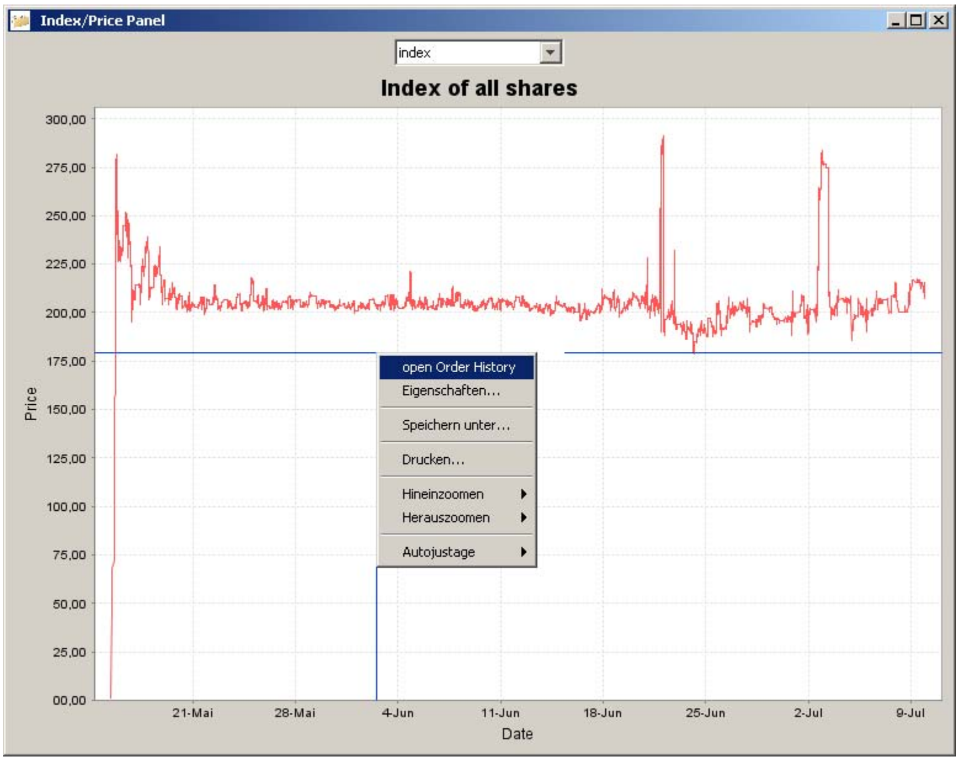

Figure 4.2: Price & index panel

4.1.3 Analyzing Price Development

The most common visualization for prices is a line plot. Prices are displayed in consecutive order and connected by lines. It is important for the user to have an easy-to-use zoom and panning functionality, a filter for single shares or combinations, and an index to observe the whole market at a glance.

The screenshot in Figure 4.2 shows the "Index/Price Panel" depicting a market index over a period of eight weeks. The context menu offers cross-links to the order history. This can be used to analyze the last arriving offers before the selected time stamp. The user can zoom in, pan, and, via the combobox above, filter for certain shares (Figure 4.3).

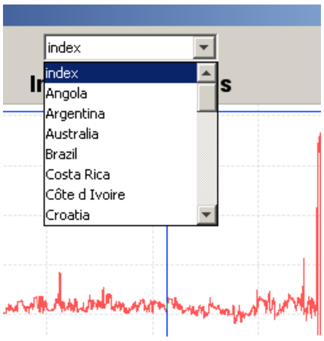

Figure 4.3: Filtering for a certain share

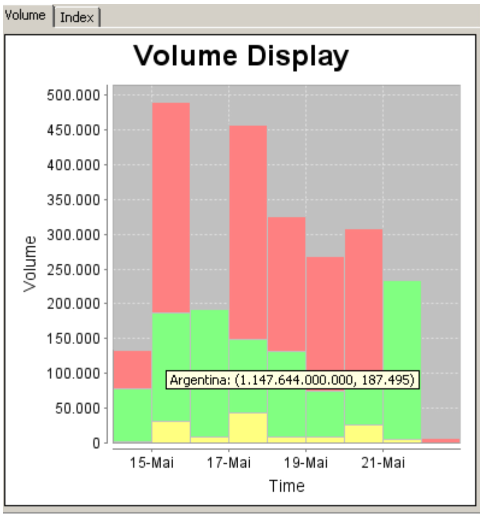

Figure 4.4: Volume panel

4.1.4 Analyzing Volume Development

Besides the price development, the overall traded volume is relevant. A stacked bar chart diagram is used to compare volumes between two different days. Each color symbolizes a different share. When moving the mouse cursor over a specific bar, the corresponding team and trading volume is displayed in a tooltip (see Figure 4.4). Similar to the *Index/Price Panel,* the user can zoom and pan in the visualization.

4.1.5 Share Inspector

The *Index/Price Panel* and the *Volume Panel* allow the user to drill down to a single share. On a share level, the user might want to better understand the order flow or transaction history at a specific point in time. Therefore, the *Share Inspector* can be used to drill down to a single transaction, order, or even an orderbook situation.

Figure 4.5 shows the *Share Inspector*. The dialog comprises three areas. On top is the toolbar where controls for share and time stamp selection are located. The time stamp can be set by entering the numbers directly, by using the pop up menu, or by using the buttons to directly change the time increment forward or back.

The second part, below the toolbar, is a tabbed view with three tabs for the different table-views (order book, orders and transactions). Each table-view has a default order criteria, which is time for orders and transactions and price followed by time for orderbooks. The user can change the ordering – indicated by the small triangles besides the column labels – by simply double clicking on the new primary order criterion. A single click on any other column label adds this label as a sub-criterion to the ordering. A second click on any criteria reverses the ordering for this criteria. Ordering is helpful to efficiently find outliers or a specific entry.

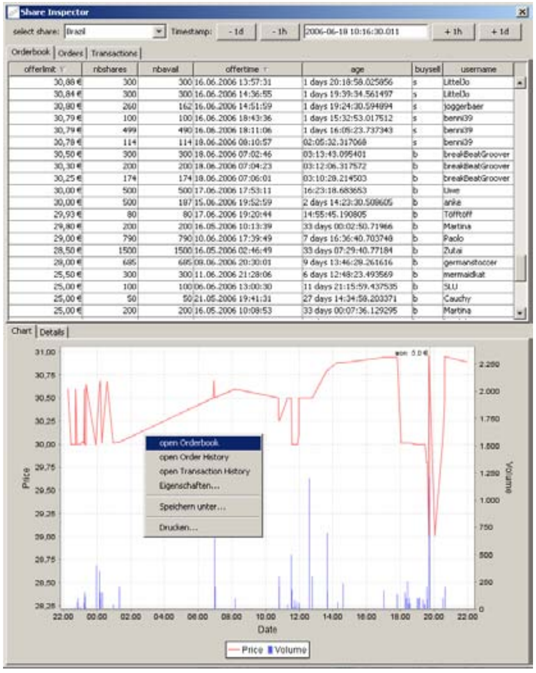

offerlimit	nbshares	nbavail	offertime	age	buysell	username
30,88 €	300	300	16.06.2006 13:57:31	1 days 20:18:58.025856	s	LittelJo
30,84 €	300	300	16.06.2006 14:36:55	1 days 19:39:34.561497	s	LittelJo
30,80 €	260	162	16.06.2006 14:51:59	1 days 19:24:30.594894	s	joggerbaer
30,79 €	100	100	16.06.2006 18:43:36	1 days 15:32:53.017512	s	benni39
30,79 €	499	490	16.06.2006 18:11:06	1 days 16:05:23.737343	s	benni39
30,78 €	114	114	18.06.2006 08:10:57	02:05:32.317068	s	benni39
30,50 €	300	300	18.06.2006 07:02:46	03:13:43.095401	b	breakBeatGroover
30,30 €	200	200	18.06.2006 07:04:23	03:12:06.317572	b	breakBeatGroover
30,25 €	174	174	18.06.2006 07:06:01	03:10:28.214503	b	breakBeatGroover
30,00 €	500	500	17.06.2006 17:53:11	16:23:18.683653	b	Uwe
30,00 €	500	187	15.06.2006 19:52:59	2 days 14:23:30.508605	b	anke
29,93 €	80	80	17.06.2006 19:20:44	14:55:45.190805	b	Tofftoff
29,80 €	200	200	16.05.2006 10:13:39	33 days 00:02:50.71966	b	Martina
29,00 €	790	790	10.06.2006 17:39:49	7 days 16:36:40.703748	b	Paolo
28,50 €	1500	1500	16.05.2006 02:46:49	33 days 07:29:40.77184	b	Zutai
28,00 €	685	685	08.06.2006 20:30:01	9 days 13:46:28.261616	b	germansoccer
25,50 €	300	300	11.06.2006 21:28:06	6 days 12:48:23.493569	b	mermaidkat
25,00 €	100	100	06.06.2006 13:00:30	11 days 21:15:59.437535	b	SLU
25,00 €	50	50	21.05.2006 19:41:31	27 days 14:34:58.203371	b	Cauchy
25,00 €	200	200	16.05.2006 10:08:53	33 days 00:07:36.129295	b	Martina

Figure 4.5: Share inspector

The third part, in the lower half of the dialog, contains general information. The *Chart*-tab plots the price and volume for a period of 24 hours centered around the current time stamp displayed in the toolbar. Besides the price and volume information, special events from the prediction market are displayed as markers in the chart. In the example in figure 4.5, Brazil played from 6 until 8 pm and eventually won. This is indicated by a gray vertical line and the text "won 5.0 €" in the chart. In the prediction market context, 5.0 € means that from then on the final value was at least five virtual currency units units (VCU). During the match, the market participants had some doubts about the outcome of the game which can be seen by the prices fluctuating between 0 and 5 VCU. The second tab contains the prediction market events and their outcomes as a table.

If the user opens the pop-up menu in the chart, the corresponding time stamp from the x-axis is used to reload the data tables. At the same time he can select whether the tab view above should switch to the order book, orders, or transactions listing. Any time stamp within the tables can be used to center the view around that time stamp via the pop up menu. All timestamps provide the pop-up menu with the option to jump to the selected time stamp. A similar function is provided for every username appearing in any of the views. If the user opens the pop-up menu or double-clicks on a username, the *Account Inspector* is opened automatically. This view is described in more detail in the following section.

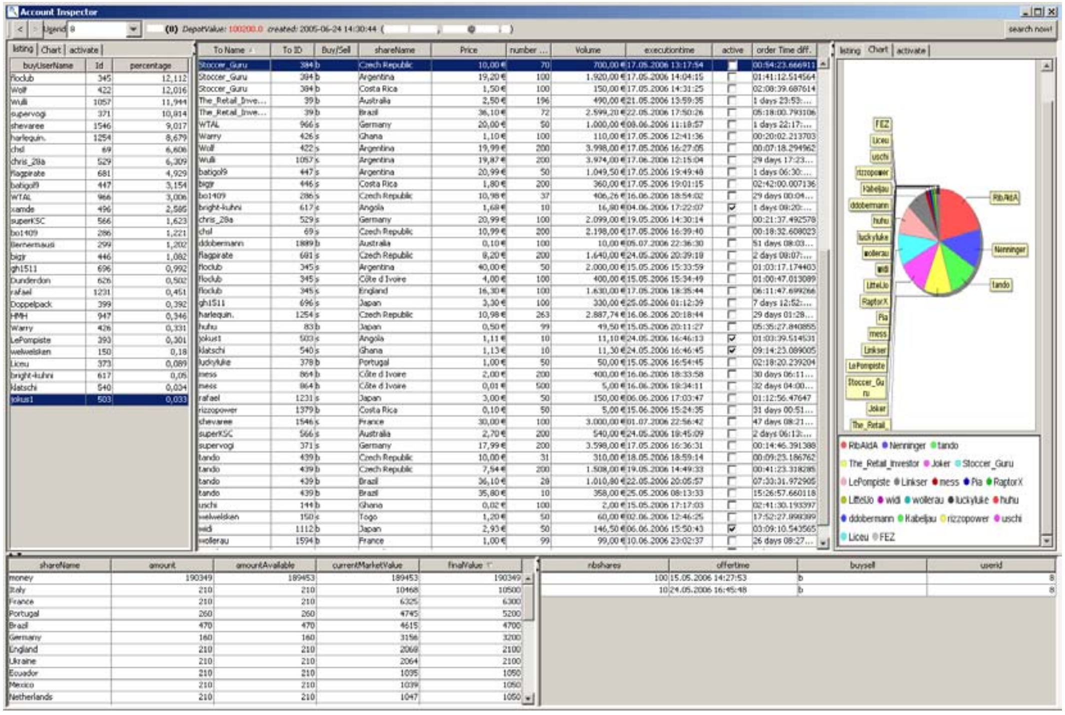

Account Inspector

< > Userid 8 (8) DepotValue: 100200.0 created: 2005-06-24 14:30:44 (,) search now!

listing | Chart | activate

buyUserName	Id	percentage
floclub	345	12,112
Wolf	422	12,016
Wulli	1057	11,944
supervogi	371	10,814
shevaree	1546	9,017
harlequin.	1254	8,679
chsl	69	6,606
chris_28a	529	6,309
flagpirate	681	4,929
batigol9	447	3,154
WTAL	966	3,006
xamde	496	2,585
superKSC	566	1,623
bo1409	286	1,221
Bernemausi	299	1,202
bigjr	446	1,082
gh1511	696	0,992
Dunderdon	626	0,502
rafael	1231	0,451
Doppelpack	399	0,392
HMH	947	0,346
Warry	426	0,331
LePompiste	393	0,301
welwelsken	150	0,18
Licnu	373	0,089
bright-kuhni	617	0,05
klatschi	540	0,034
jokus1	503	0,033

To Name	To ID	Buy/Sell	shareName	Price	number ...	Volume	executiontime	active	order Time diff.
Stoccer_Guru	384	b	Czech Republic	10,00 €	70	700,00 €	17.05.2006 13:17:54	□	00:54:23.666911
Stoccer_Guru	384	b	Argentina	19,20 €	100	1.920,00 €	17.05.2006 14:04:15	□	01:41:12.514564
Stoccer_Guru	384	b	Costa Rica	1,50 €	100	150,00 €	17.05.2006 14:31:25	□	02:08:39.687614
The_Retail_Inve...	39	b	Australia	2,50 €	196	490,00 €	21.05.2006 13:59:35	□	1 days 23:53:...
The_Retail_Inve...	39	b	Brazil	36,10 €	72	2.599,20 €	22.05.2006 17:50:26	□	05:18:00.793106
WTAL	966	s	Germany	20,00 €	50	1.000,00 €	08.06.2006 11:18:57	□	1 days 22:17:...
Warry	426	s	Ghana	1,10 €	100	110,00 €	17.05.2006 12:41:36	□	00:20:02.213703
Wolf	422	s	Argentina	19,99 €	200	3.998,00 €	17.05.2006 16:27:05	□	00:07:18.294962
Wulli	1057	s	Argentina	19,87 €	200	3.974,00 €	17.06.2006 12:15:04	□	29 days 17:23...
batigol9	447	s	Argentina	20,99 €	50	1.049,50 €	17.05.2006 19:49:48	□	1 days 06:30:...
bigjr	446	s	Costa Rica	1,80 €	200	360,00 €	17.05.2006 19:01:15	□	02:42:00.007136
bo1409	286	s	Czech Republic	10,98 €	37	406,26 €	16.06.2006 18:54:02	□	29 days 00:04...
bright-kuhni	617	s	Angola	1,68 €	10	16,80 €	04.06.2006 17:22:07	☑	1 days 08:20:...
chris_28a	529	s	Germany	20,99 €	100	2.099,00 €	19.05.2006 14:30:14	□	00:21:37.492578
chsl	69	s	Czech Republic	10,99 €	200	2.198,00 €	17.05.2006 16:39:40	□	00:18:32.608023
ddobermann	1889	b	Australia	0,10 €	100	10,00 €	05.07.2006 22:36:30	□	51 days 08:03...
flagpirate	681	s	Czech Republic	8,20 €	200	1.640,00 €	24.05.2006 20:39:18	□	2 days 08:07:...
floclub	345	s	Argentina	40,00 €	50	2.000,00 €	15.05.2006 15:33:59	□	01:03:17.174403
floclub	345	s	Côte d Ivoire	4,00 €	100	400,00 €	15.05.2006 15:34:49	□	01:00:47.013089
floclub	345	s	England	16,30 €	100	1.630,00 €	17.05.2006 18:35:44	□	06:11:47.699266
gh1511	696	s	Japan	3,30 €	100	330,00 €	25.05.2006 01:12:39	□	7 days 12:52:...
harlequin.	1254	s	Czech Republic	10,98 €	263	2.887,74 €	16.06.2006 20:18:44	□	29 days 01:28...
huhu	83	b	Japan	0,50 €	99	49,50 €	15.05.2006 20:11:27	□	05:35:27.840855
jokus1	503	s	Angola	1,11 €	10	11,10 €	24.05.2006 16:46:13	☑	01:03:39.514531
klatschi	540	s	Ghana	1,13 €	10	11,30 €	24.05.2006 16:46:45	☑	09:14:23.089005
luckyluke	378	b	Portugal	1,00 €	50	50,00 €	15.05.2006 16:54:45	□	02:18:20.239204
mess	864	b	Côte d Ivoire	2,00 €	200	400,00 €	16.06.2006 18:33:58	□	30 days 06:11...
mess	864	b	Côte d Ivoire	0,01 €	500	5,00 €	16.06.2006 18:34:11	□	32 days 04:00...
rafael	1231	s	Japan	3,00 €	50	150,00 €	06.06.2006 17:03:47	□	01:12:56.47647
rizzopower	1379	b	Costa Rica	0,10 €	50	5,00 €	15.06.2006 15:24:35	□	31 days 00:51...
shevaree	1546	s	France	30,00 €	100	3.000,00 €	01.07.2006 22:56:42	□	47 days 08:21...
superKSC	566	s	Australia	2,70 €	200	540,00 €	24.05.2006 18:45:09	□	2 days 06:13:...
supervogi	371	s	Germany	17,99 €	200	3.598,00 €	17.05.2006 16:36:31	□	00:14:46.391388
tando	439	b	Czech Republic	10,00 €	31	310,00 €	18.05.2006 18:59:14	□	00:09:23.186762
tando	439	b	Czech Republic	7,54 €	200	1.508,00 €	19.05.2006 14:49:33	□	00:41:23.318285
tando	439	b	Brazil	36,10 €	28	1.010,80 €	22.05.2006 20:05:57	□	07:33:31.972905
tando	439	b	Brazil	35,80 €	10	358,00 €	25.05.2006 08:13:33	□	15:26:57.660118
uschi	144	b	Ghana	0,02 €	100	2,00 €	15.05.2006 17:17:03	□	02:41:30.193397
welwelsken	150	s	Togo	1,20 €	50	60,00 €	02.06.2006 12:46:25	□	17:52:27.898389
widi	1112	b	Japan	2,93 €	50	146,50 €	06.06.2006 15:50:43	☑	03:09:10.543565
wollerau	1594	b	France	1,00 €	99	99,00 €	10.06.2006 23:02:37	□	26 days 08:27...

listing | Chart | activate

FEZ
Liceu
uschi
rizzopower
Kabeljau
ddobermann
huhu
luckyluke
wollerau
widi
LittleJo
RaptorX
Pia
mess
Linkser
LePompiste
Stoccer_Guru
Joker
The_Retail_
RibAldA
Nenninger
tando
● RibAldA ● Nenninger ● tando
● The_Retail_Investor ● Joker ● Stoccer_Guru
● LePompiste ● Linkser ● mess ● Pia ● RaptorX
● LittleJo ● widi ● wollerau ● luckyluke ● huhu
● ddobermann ● Kabeljau ● rizzopower ● uschi
● Liceu ● FEZ

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	190349	189453	189453	190349
Italy	210	210	10468	10500
France	210	210	6325	6300
Portugal	260	260	4745	5200
Brazil	470	470	4615	4700
Germany	160	160	3156	3200
England	210	210	2069	2100
Ukraine	210	210	2064	2100
Ecuador	210	210	1035	1050
Mexico	210	210	1039	1050
Netherlands	210	210	1047	1050

nbshares	offertime	buysell	userid
100	15.05.2006 14:27:53	b	8
10	24.05.2006 16:45:48	b	8

Figure 4.6: Account inspector

4.1.6 Account Inspector

According to the paradigm that markets are formed by social entities, accounts play a key role. Some entities have to be investigated in more detail if they have been reported by an algorithm or gotten the attention of a market operator. All relevant account information is provided by the *Account Inspector*. This dialog offers views for every activity an account may have conducted (see Figure 4.6).

The dialog is comprisesd of five parts. A tool bar at the top allows the user to enter a specific username or id, step backward and forward in the history, or directly select a previously investigated account via the combobox. User registration details (account name, id, current portfolio value, creation time stamp, forename, surname and email) are displayed next to the combobox. The personal information is anonymized in the screenshot. The red color of the portfolio value indicates that the account lost money.

The center part of the *Account Inspector* is split into three columns. In the second column, all transactions of the account are listed. Besides the information about id and name of buyer and seller, amount, price, volume, and time stamp, the investigator can see whether the currently investigated account actively matched against an existing order in the order book – indicated by the column *active* – and the delay between the two matching orders. For instance, in the first line, a user's buy order for *Czech Republic* was actively matched by *STOCCER_Guru* with a sell order of over 70 shares for 700.00 on May, 17th at 13:17:54. The buyer had put in his order 54 minutes and 23 seconds before. This will be the main column the investigator is focusing on.

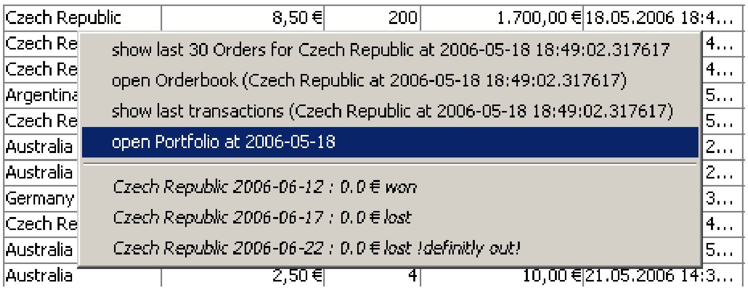
Czech Republic | 8,50 € | 200 | 1.700,00 € | 18.05.2006 18:4...
show last 30 Orders for Czech Republic at 2006-05-18 18:49:02.317617
open Orderbook (Czech Republic at 2006-05-18 18:49:02.317617)
show last transactions (Czech Republic at 2006-05-18 18:49:02.317617)
open Portfolio at 2006-05-18
Czech Republic 2006-06-12 : 0.0 € won
Czech Republic 2006-06-17 : 0.0 € lost
Czech Republic 2006-06-22 : 0.0 € lost !definitly out!
Czech Re
Czech Re
Argentina
Czech Re
Australia
Australia
Germany
Czech Re
Australia
Australia | 2,50 € | 4 | 10,00 € | 21.05.2006 14:3...

Figure 4.7: Pop-up menu

This listing is framed by lists of all buying and selling partners. Each view can be shown as either a table or as a pie chart diagram. In Figure 4.6, the accounts who bought shares from the account currently being investigated are shown in the table format on the left side; while on the right, all accounts who sold shares to the currently investigated account are displayed as a pie chart. The order criteria for both sides (sell and buy) is the volume, which is calculated as price multiplied by the number of shares. An account is ranked highest if it bought the highest volume from the current account. Double clicking on any of the listed accounts in the selling-, buying- or transaction partner panels switches the selected account into the currently investigated account. If the user wants to go back to the previous account, he can use the browser navigation elements of the toolbar.

The bottom of the *Account Inspector* is split into two columns. In the first column, portfolio information is depicted. By default, the current portfolio situation is shown in descending order of the current market value. For each share in the property of the account, the number of shares bought so far (*amount*), the number of shares which are not covered by open sell orders (amount available), the value at current market prices, and the value multiplied by the final payoff for each share is listed. If the final payoff is unknown, this column is zero. Trades are listed in the right table portfolio. For detailed information about portfolio trading refer to Section 1.2 on page 4.

If the user investigates a transaction, the price or quantity can be judged better in the context of the price development of the share and the portfolio the user had at that point in time. The pop-up menu (see Figure 4.7) in the share column of the transaction table allows the user to access the *Share Inspector*, to open a portfolio view, and list all scheduled events (in the STOCCER context soccer matches) for this share. In the lower half of the menu, all matches of *Czech Republic* are displayed. The team played three matches, but they were all after May 18th and are therefore displayed in italics. Past matches are displayed in bold. This way the user exactly knows which information is most important about this share. Since the screenshot was taken after the soccer world championships had finished, the results of all matches were known and printed next to each match. Since *Czech Republic* lost two of its three matches, they did not enter the second half of the championships and the value of this share dropped to zero.

A separate dialog can be opened to see the portfolio of the trader at the moment of the

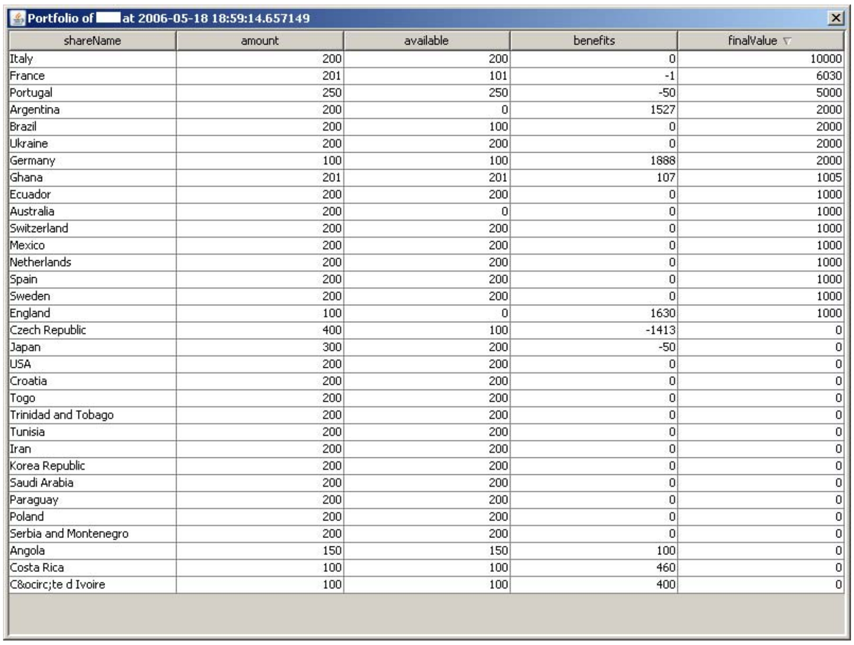

Portfolio of [redacted] at 2006-05-18 18:59:14.657149

shareName	amount	available	benefits	finalValue ▽
Italy	200	200	0	10000
France	201	101	-1	6030
Portugal	250	250	-50	5000
Argentina	200	0	1527	2000
Brazil	200	100	0	2000
Ukraine	200	200	0	2000
Germany	100	100	1888	2000
Ghana	201	201	107	1005
Ecuador	200	200	0	1000
Australia	200	0	0	1000
Switzerland	200	200	0	1000
Mexico	200	200	0	1000
Netherlands	200	200	0	1000
Spain	200	200	0	1000
Sweden	200	200	0	1000
England	100	0	1630	1000
Czech Republic	400	100	-1413	0
Japan	300	200	-50	0
USA	200	200	0	0
Croatia	200	200	0	0
Togo	200	200	0	0
Trinidad and Tobago	200	200	0	0
Tunisia	200	200	0	0
Iran	200	200	0	0
Korea Republic	200	200	0	0
Saudi Arabia	200	200	0	0
Paraguay	200	200	0	0
Poland	200	200	0	0
Serbia and Montenegro	200	200	0	0
Angola	150	150	100	0
Costa Rica	100	100	460	0
C&ocirc;te d Ivoire	100	100	400	0

Figure 4.8: Portfolio dialog

transaction (see Figure 4.8). The table is organized similarly to the final portfolio view in the *Account Inspector*. An additional column *benefits* indicates whether the account holder had gains or losses from this share. The *benefit* is the sum of all selling minus the sum of all buying transactions. Portfolio trades are not considered because, on the one hand, most people do not trade whole portfolios, on the other hand, this simplified formula allows the investigator to distinguish where the user had considerable gains or losses and which shares he traded actively. Though exact figures may differ with portfolio trades included, the proportion is similar in most cases.

4.2 Combining Algorithms and Visualization

4.2.1 Active Search Support

The integration of the fraud detection algorithms from Chapter 2 is realized via two mechanisms. If the user wants to actively search for accounts, he can start and configure the algorithms via the menu from the main application window. Figure 4.9 shows a screenshot of the algorithm selection dialog. Several algorithms can be selected and configured individually. To configure an algorithm, the user selects the *Properties* button beside the selected algorithm. An example of the properties dialog for the Gini-algorithm is given in Figure 4.10 (a). If the user is new to the market, he may select the preview button to see the scoring distribution of the examined accounts using his current settings (c.f. Figure 4.10 (b)). After the algorithms have been started, they are sequentially executed in the background, so the user can continue working with the application, e.g.

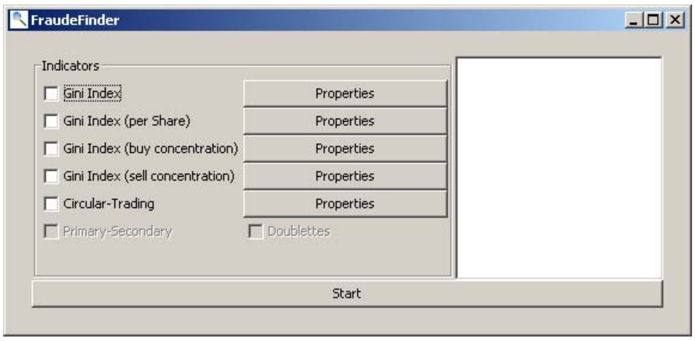

FraudeFinder

Indicators

□ Gini Index

□ Gini Index (per Share)

□ Gini Index (buy concentration)

□ Gini Index (sell concentration)

□ Circular-Trading

□ Primary-Secondary

□ Doublettes

Figure 4.9: Algorithm selection

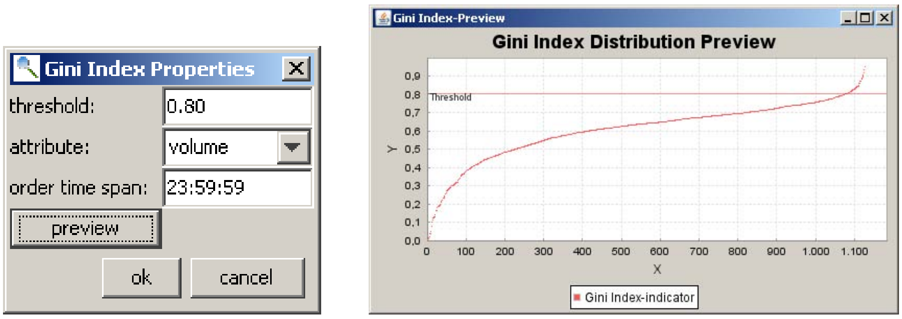

Gini Index Properties

threshold: 0.80

attribute: volume

order time span: 23:59:59

(a) properties (b) distribution preview

Figure 4.10: Configuration of the algorithm properties

investigating further accounts or transactions. The algorithm results will pop up in a results window after the last algorithm has finished. Additionally, the user is informed by an acoustic signal that his search has finished.

Cases are listed one per row. Each case consists of the two partner accounts, the indicator which reported the case, the probability ranking assigned by the indicator, the involved volume (price * number of shares), the number of transactions, and the number of shares. A toolbar above the search results table (see Figure 4.11) provides controls to filter the list of reported cases in several advanced ways:

- The text field provides an incremental search for the entered character sequence. This feature can be used to see how often, with which partners, and which indicators an account has been reported. In fact, some people labeled the usernames of their different accounts very similarly; thus the whole group appeared when investigating the first suspicious account.
- The second button (“add selection to *CaseManager*”) is used by the investigator to confirm one or more cases. These cases are inserted into the central case database which is called *CaseManager*. Each account referenced by a case in the *CaseManager* will be highlighted in the graph visualization. Thus, the monitoring staff is alerted whenever the concluding node is a red one instead of a blue one. All cases are listed in the context menu of the node, in order to allow fast user access to this information.

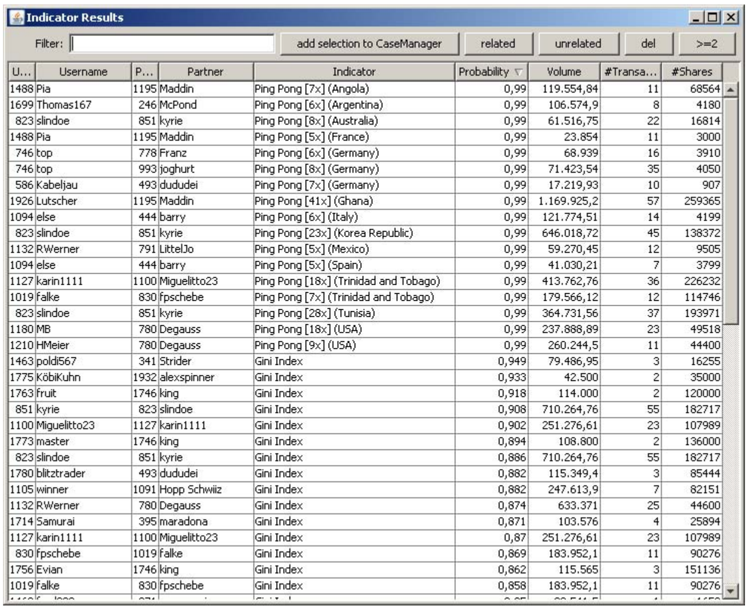

U...	Username	P...	Partner	Indicator	Probability ▽	Volume	#Transa...	#Shares
1488	Pia	1195	Maddin	Ping Pong [7x] (Angola)	0,99	119.554,84	11	68564
1699	Thomas167	246	McPond	Ping Pong [6x] (Argentina)	0,99	106.574,9	8	4180
823	slindoe	851	kyrie	Ping Pong [8x] (Australia)	0,99	61.516,75	22	16814
1488	Pia	1195	Maddin	Ping Pong [5x] (France)	0,99	23.854	11	3000
746	top	778	Franz	Ping Pong [6x] (Germany)	0,99	68.939	16	3910
746	top	993	joghurt	Ping Pong [8x] (Germany)	0,99	71.423,54	35	4050
586	Kabeljau	493	dududei	Ping Pong [7x] (Germany)	0,99	17.219,93	10	907
1926	Lutscher	1195	Maddin	Ping Pong [41x] (Ghana)	0,99	1.169.925,2	57	259365
1094	else	444	barry	Ping Pong [6x] (Italy)	0,99	121.774,51	14	4199
823	slindoe	851	kyrie	Ping Pong [23x] (Korea Republic)	0,99	646.018,72	45	138372
1132	RWerner	791	LittelJo	Ping Pong [5x] (Mexico)	0,99	59.270,45	12	9505
1094	else	444	barry	Ping Pong [5x] (Spain)	0,99	41.030,21	7	3799
1127	karin1111	1100	Miguelitto23	Ping Pong [18x] (Trinidad and Tobago)	0,99	413.762,76	36	226232
1019	falke	830	fpschebe	Ping Pong [7x] (Trinidad and Tobago)	0,99	179.566,12	12	114746
823	slindoe	851	kyrie	Ping Pong [28x] (Tunisia)	0,99	364.731,56	37	193971
1180	MB	780	Degauss	Ping Pong [18x] (USA)	0,99	237.888,89	23	49518
1210	HMeier	780	Degauss	Ping Pong [9x] (USA)	0,99	260.244,5	11	44400
1463	poldi567	341	Strider	Gini Index	0,949	79.486,95	3	16255
1775	KöbiKuhn	1932	alexspinner	Gini Index	0,933	42.500	2	35000
1763	fruit	1746	king	Gini Index	0,918	114.000	2	120000
851	kyrie	823	slindoe	Gini Index	0,908	710.264,76	55	182717
1100	Miguelitto23	1127	karin1111	Gini Index	0,902	251.276,61	23	107989
1773	master	1746	king	Gini Index	0,894	108.800	2	136000
823	slindoe	851	kyrie	Gini Index	0,886	710.264,76	55	182717
1780	blitztrader	493	dududei	Gini Index	0,882	115.349,4	3	85444
1105	winner	1091	Hopp Schwiiz	Gini Index	0,882	247.613,9	7	82151
1132	RWerner	780	Degauss	Gini Index	0,874	633.371	25	44600
1714	Samurai	395	maradona	Gini Index	0,871	103.576	4	25894
1127	karin1111	1100	Miguelitto23	Gini Index	0,87	251.276,61	23	107989
830	fpschebe	1019	falke	Gini Index	0,869	183.952,1	11	90276
1756	Evian	1746	king	Gini Index	0,862	115.565	3	151136
1019	falke	830	fpschebe	Gini Index	0,858	183.952,1	11	90276

Figure 4.11: Indicator result table

- After running the application for a while, the user may have already run the search or may have tagged suspicious accounts manually. If the user now wants to find further evidence for his suspicion, he might be interested in finding further indications of fraud for this set of accounts. This is supported by the *related* button. When pushed, the application removes all cases from the list where neither of the two accounts already appears in at least one case of the *CaseManager*.
- On the other hand, the *unrelated* button removes all cases where at least one of the two accounts already appears in at least one case of the *CaseManager*. This supports the user when he is only interested in new cases he has not confirmed so far.
- If the user is sure that a case is irrelevant, he may delete the case from the list using the *del* button. A future enhancement would be to use this information as training or calibration information for the algorithms.
- The *>=2* button allows the user to only filter for cases where the account combination is reported at least twice. The filter allows inverted couples as well. This means that "a,b" is treated the same as "b,a". This option is especially useful if several algorithms were executed sequentially. If the user wants to check which couples were reported from more than one algorithm, this option assists him in finding them.

The *CaseManager* can be opened directly via the menu bar of the application window. The appearance of the *CaseManager* is very similar to the indicator results dialog. The main difference is that the *CaseManager* allows the user to import and export the central case database to a file. The database file is a comma separated value list that can be

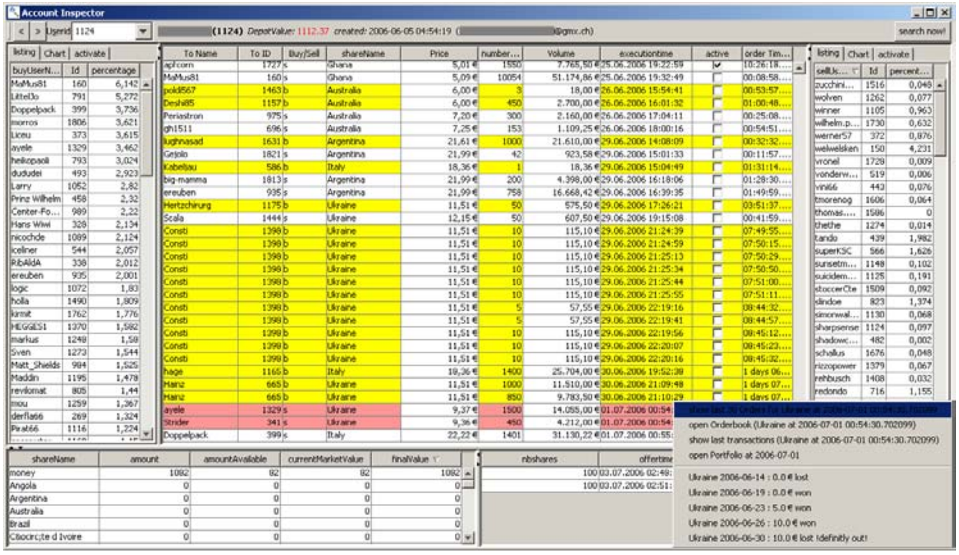

Account Inspector

Userid 1124

[redacted](1124) DepotValue: 1112.37 created: 2006-06-05 04:54:19 ([redacted]@gmx.ch)

search now!

listing | Chart | activate

buyUserN...	Id	percentage
MaMus81	160	6,142
LittelJo	791	5,272
Doppelpack	399	3,736
morros	1806	3,621
Liceu	373	3,615
ayele	1329	3,462
heikopaoli	793	3,024
dududei	493	2,923
Larry	1052	2,82
Prinz Wilhelm	458	2,32
Center-Fo...	989	2,22
Hans Wiwi	328	2,134
nicochde	1089	2,124
icelner	544	2,057
RibAldA	338	2,012
ereuben	935	2,001
logic	1072	1,83
holla	1490	1,809
kirmit	1762	1,776
HEGGES1	1370	1,582
markus	1248	1,58
Sven	1273	1,544
Matt_Shields	984	1,525
Maddin	1195	1,478
revilomat	805	1,44
mou	1259	1,367
derfla66	269	1,324
Pirat66	1116	1,224

To Name	To ID	Buy/Sell	shareName	Price	number...	Volume	executiontime	active	order Tim...
apfcorn	1727	s	Ghana	5,01 €	1550	7.765,50 €	25.06.2006 19:22:59	☑	10:26:18....
MaMus81	160	s	Ghana	5,09 €	10054	51.174,86 €	25.06.2006 19:32:49	□	00:08:58....
poldi567	1463	b	Australia	6,00 €	3	18,00 €	26.06.2006 15:54:41	□	00:53:57....
Deshi85	1157	b	Australia	6,00 €	450	2.700,00 €	26.06.2006 16:01:32	□	01:00:48....
Periastron	975	s	Australia	7,20 €	300	2.160,00 €	26.06.2006 17:04:11	□	00:25:08....
gh1511	696	s	Australia	7,25 €	153	1.109,25 €	26.06.2006 18:00:16	□	00:54:51....
lughnasad	1631	b	Argentina	21,61 €	1000	21.610,00 €	29.06.2006 14:08:09	□	00:32:32....
Gejolo	1821	s	Argentina	21,99 €	42	923,58 €	29.06.2006 15:01:33	□	00:11:57....
Kabeljau	586	b	Italy	18,36 €	1	18,36 €	29.06.2006 15:04:49	□	01:31:14....
big-mamma	1813	s	Argentina	21,99 €	200	4.398,00 €	29.06.2006 16:18:06	□	01:28:30....
ereuben	935	s	Argentina	21,99 €	758	16.668,42 €	29.06.2006 16:39:35	□	01:49:59....
Hertzchirurg	1175	b	Ukraine	11,51 €	50	575,50 €	29.06.2006 17:26:21	□	03:51:37....
Scala	1444	s	Ukraine	12,15 €	50	607,50 €	29.06.2006 19:15:08	□	00:41:59....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 21:24:39	□	07:49:55....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 21:24:59	□	07:50:15....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 21:25:13	□	07:50:29....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 21:25:34	□	07:50:50....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 21:25:44	□	07:51:00....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 21:25:55	□	07:51:11....
Consti	1398	b	Ukraine	11,51 €	5	57,55 €	29.06.2006 22:19:16	□	08:44:32....
Consti	1398	b	Ukraine	11,51 €	5	57,55 €	29.06.2006 22:19:41	□	08:44:57....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 22:19:56	□	08:45:12....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 22:20:07	□	08:45:23....
Consti	1398	b	Ukraine	11,51 €	10	115,10 €	29.06.2006 22:20:16	□	08:45:32....
hage	1165	b	Italy	18,36 €	1400	25.704,00 €	30.06.2006 19:52:38	□	1 days 06...
Mainz	665	b	Ukraine	11,51 €	1000	11.510,00 €	30.06.2006 21:09:48	□	1 days 07...
Mainz	665	b	Ukraine	11,51 €	850	9.783,50 €	30.06.2006 21:10:29	□	1 days 07...
ayele	1329	s	Ukraine	9,37 €	1500	14.055,00 €	01.07.2006 00:54		
Strider	341	s	Ukraine	9,36 €	450	4.212,00 €	01.07.2006 00:54		
Doppelpack	399	s	Italy	22,22 €	1401	31.130,22 €	01.07.2006 00:55:		

listing | Chart | activate

sellUs...	Id	percent...
zucchini...	1516	0,048
wolven	1262	0,077
winner	1105	0,963
wilhelm.p...	1730	0,632
werner57	372	0,876
welwelsken	150	4,231
vronel	1728	0,009
vonderw...	519	0,006
vini66	443	0,076
tmorenog	1606	0,064
thomas....	1586	0
thethe	1274	0,014
tando	439	1,982
superKSC	566	1,626
sunsetm...	1148	0,102
suicidem...	1125	0,191
stoccerCte	1509	0,092
dindoe	823	1,374
simonwal...	1130	0,068
sharpsense	1124	0,097
shadowc...	482	0,002
schallus	1676	0,048
rizzopower	1379	0,067
rehbusch	1408	0,032
redondo	716	1,155

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	1082	82	82	1082
Angola	0	0	0	0
Argentina	0	0	0	0
Australia	0	0	0	0
Brazil	0	0	0	0
C&ocirc;te d Ivoire	0	0	0	0

nbshares	offertime
100	03.07.2006 02:48:
100	03.07.2006 02:51:

show last 30 Orders for Ukraine at 2006-07-01 00:54:30.702099
open Orderbook (Ukraine at 2006-07-01 00:54:30.702099)
show last transactions (Ukraine at 2006-07-01 00:54:30.702099)
open Portfolio at 2006-07-01

Ukraine 2006-06-14 : 0.0 € lost
Ukraine 2006-06-19 : 0.0 € won
Ukraine 2006-06-23 : 5.0 € won
Ukraine 2006-06-26 : 10.0 € won
Ukraine 2006-06-30 : 10.0 € lost !definitly out!

Figure 4.12: Transaction highlighting

opened in Excel and many other applications. This feature allows the preservation of the investigation results throughout several sessions with the prototype.

4.2.2 Indirect Active Search Support

Besides the active search, the algorithms can also evaluate transactions in the background and support the user by highlighting suspicious transactions within the table views. Different colors can be used to distinguish different reasons for suspicion. Figure 4.12 shows an example screenshot. In the screenshot, two of the four signal colors are depicted in the *Ukraine* transactions: the sell transaction to *Scala* on the 29^{th} is colored white (the default color) since it is not considered suspicious but rather of natural interest for the seller to sell over price. The remaining buy transactions are colored yellow to inform the user that this participant bought the share above the current assured market price. Again, this is common habit on financial and prediction markets, usually in the hope of rising prices. In this case, it is not an unrealistic expectation since *Ukraine* had won 3:0 against *Switzerland* on the 26^{th}. The last two transactions are colored in red because the participant is selling *Ukraine* shares below the definite market value (see the context menu information on the right). If he kept them till the end of the market he would have gained a difference of:

$$1950 * 10.00 - \ (1500 * 9.37 + 450 * 9.36) = 1233.00$$

However, this does not necessarily have to be fraud. It may have been a mistake, ignorance, or an attempt to gain liquidity for another planned investment into a different share. The user's task is to evaluate these possibilities. He may check the price development of *Ukraine* after May 30^{th} with the *Share Inspector* to see whether it was just

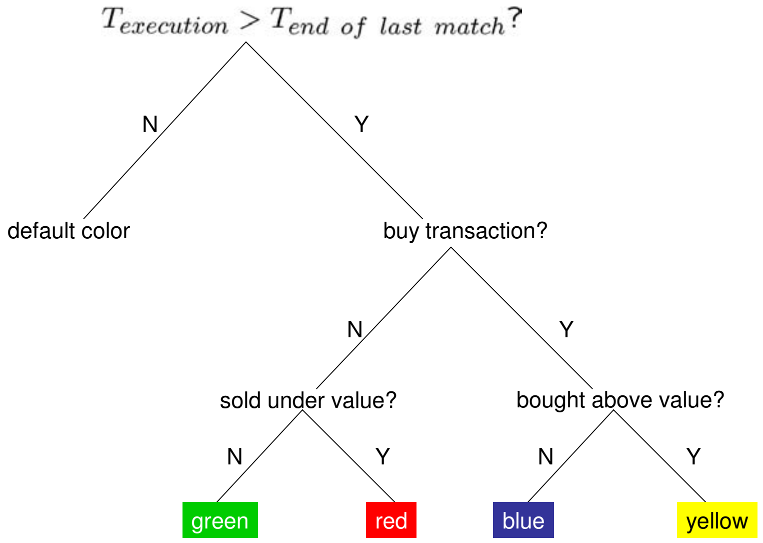

Figure 4.13: Transaction coloring decision tree

ignorance, to look for similar registration data between the involved accounts, to look into whether they helped each other, and to see if the transaction was completed to use the liquidity for another investment. The colors are selected according to the priority of the case. If someone is selling a share above its common value, the transaction is marked green to point out the benefits. A transaction is marked in blue if the buyer got the share below the common (minimum) value. The decision tree to determine the color for a transaction is shown in Figure 4.13. Similar mechanisms exist for the graph visualization from the Chapter 3.

4.2.3 Monitoring Support

The main application window combines both the graph display and traditional charts. In Figure 4.14, the graph display can be seen in the main window. On the right are three smaller displays: the upper one represents a map view of the graph, the middle one represents a trading volume bar chart, and the bottom one represents the price and index chart. The toolbar on top contains the controls to set and change the time stamp and stepping, the share filter and the sliding window. The share filter allows for filtering the visualizations in all displays for transactions only resulting from a selected non-empty subset of shares. This enables the user to relate the activity observed in the market displays (volume and price) to the activity of the graph.

Besides the share filtering, the most important complexity reduction stems from the sliding window. The settings for the sliding window can be configured using the controls

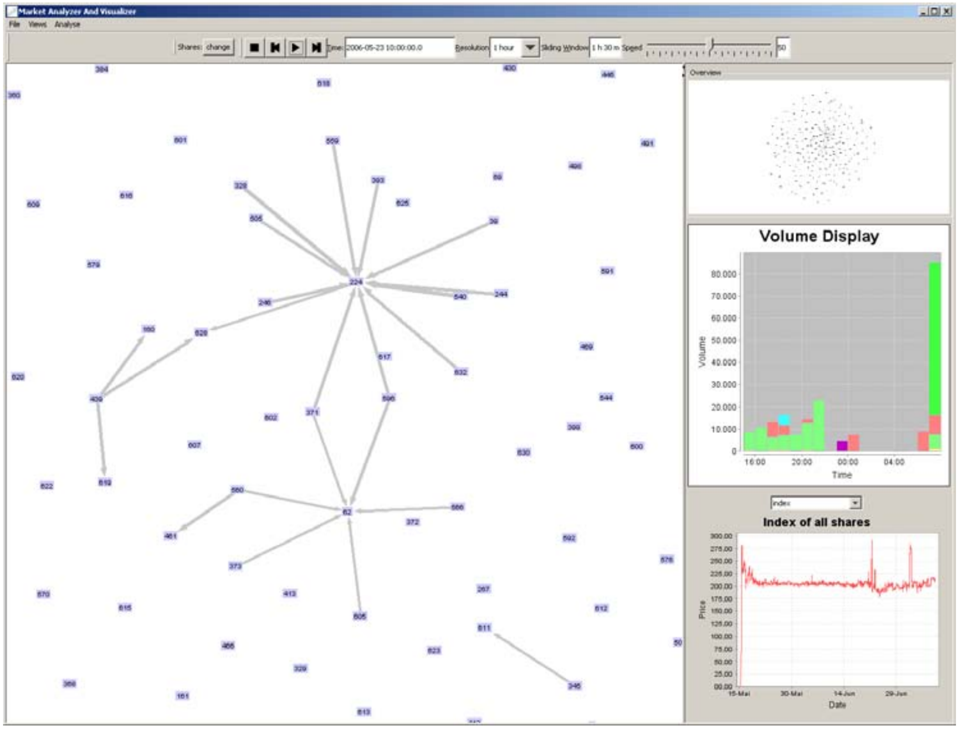

Market Analyzer And Visualizer

Shares: change

Time: 2006-05-23 10:00:00.0

Resolution: 1 hour

Sliding Window: 1 h 30 m

Speed: 50

Figure 4.14: Screenshot of the market replay

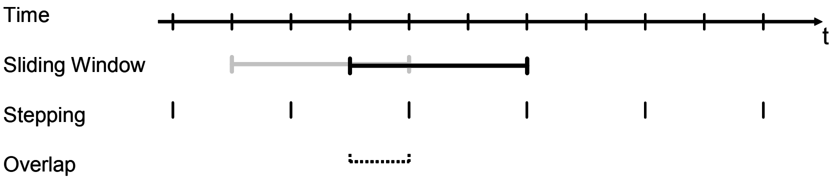

Figure 4.15: Sliding window

on the toolbar. On the left, the user can start, pause, and stop stepping through the data set. The step size is determined by the *Resolution* combobox. Selecting either the step forward or backward buttons allows the user to progress one step at a time, while pressing the play button moves on to the next step automatically after a certain delay. The delay can be configured via the speed slider on the right. The sliding window length, which controls the resolution of the step on the time axis, can be configured via the sliding window text field. Figure 4.15 helps to explain the difference between the stepping and sliding window. The time axis is divided into equal units (e.g. hours in the sample). This means that the visualization is depicting the graph at 2 am, 4 am, and so on. The sliding window sets the duration of the trading period the graph edges reflect. In Figure 4.15, the sliding window is set to three units, which in the sample equals three hours. This creates an overlap of one hour, which is half of the length of the stepping.

Though it seems redundant at first sight, since we repeat information already presented to the user, it helps the user to maintain his mental map in a quickly changing visualization. Since some edges of the former situation are repeated, several nodes stay in the same

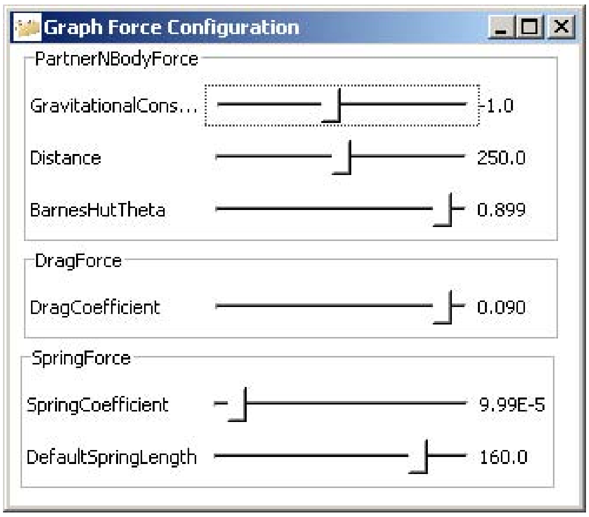
Graph Force Configuration

PartnerNBodyForce

GravitationalCons... 1.0

Distance 250.0

BarnesHutTheta 0.899

DragForce

DragCoefficient 0.090

SpringForce

SpringCoefficient 9.99E-5

DefaultSpringLength 160.0

Figure 4.16: Layout parameter dialog

or similar location. This reduces the effort the user's short time memory has to spend following the activity in the market. The transition states of the visualization are animated by the prefuse toolkit. Thus the user gets a smooth animation from the layout of the previous step to the current step.

Though it is usually not necessary, the user may adjust the layout parameters of the graph visualization according to his preferences. Figure 4.16 shows a screenshot of the dialog. Most of the parameters concern force functions of the basic algorithm. The first block configures the gravitation constant c_g of the repelling force function (see Section 3.1) and two other parameters of the Barnes-Hut-extension. The second block configures the drag force which influences how fast nodes move over the background. Finally, the last block adjusts the spring coefficient c_δ and spring length l of the spring forces (see Section 3.2). Manipulating the parameters is reflected in the visualization in real time. This allows the user to efficiently adjust the visualization.

4.3 User Interaction

Interaction is a key feature in visual analytics and also an important requirement for our visualization. The graph view offers a high number of interaction possibilities. The mouse scroll wheel controls the zoom level of the display. By left clicking on the background (or an edge) and dragging in any direction, the user can pan (move the visualization) to see other areas or center the clipping area on a new position. By left clicking on a node and dragging, the user can drag the node to a new position. This can be used to assist the layouter or give the visualization a special ordering. After dropping the node on a new position, the layout algorithm will start optimizing the layout from the new starting positions. To keep such a manual layout, the user can freeze all nodes in their current positions. The pop-up menu (see Figure 4.17 (a)) contains two options to freeze and

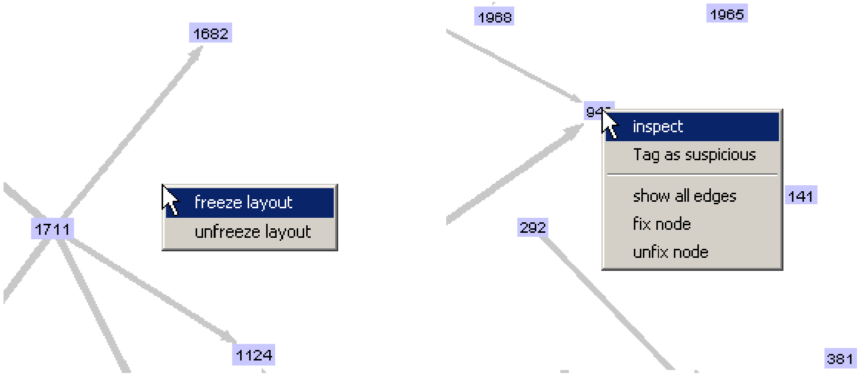

(a) Background pop-up menu (b) Node pop-up menu

Figure 4.17: Pop-up menus of the graph display

unfreeze the layout.

If a node has to be observed in more detail, the user can pin it on the background by selecting the *fix node* option from the pop up menu (see Figure 4.17 (b)) of the node. If desired, several nodes can be fixed at the same time. This way the user can arrange nodes at more prominent spots. If the user is interested in all former trading partners, he can deactivate the sliding window time filter for this node from the pop-up menu. This will display filtered and unfiltered edges of the selected node on the screen, arranging all trading partners around the selected node. Nodes can be tagged to keep track of a suspicious account. After tagging a node, its color changes to red and the text of the tag is displayed in the pop up menu.

Through a context menu available on edges, nodes, bar charts and scatter plots, the user can drill down to the account and transactions represented by the selected element. Since a very common use case is to investigate around an interesting edge or node, a double click on a node will open the *Account Inspector* and load the account. If an edge is under investigation (e.g. by right clicking on it), the *Account Inspector* will load the source node and highlight all transactions to the destination node. This will automatically scroll to the visible part of the table enabling the user to drill down in a quick and efficient way.

4.4 Further Dialogs

4.4.1 Market Selection

The first dialog that appears when the prototype is started is the *market selection dialog* (see Figure 4.19). In this dialog, the user can not only select the market he wants to monitor or investigate, but he can also get some information about the market. Besides the market name, the first day the market opened and the last day it will close is shown. Lastly, the dialog lists the current number of orders and the number of concluded transactions.

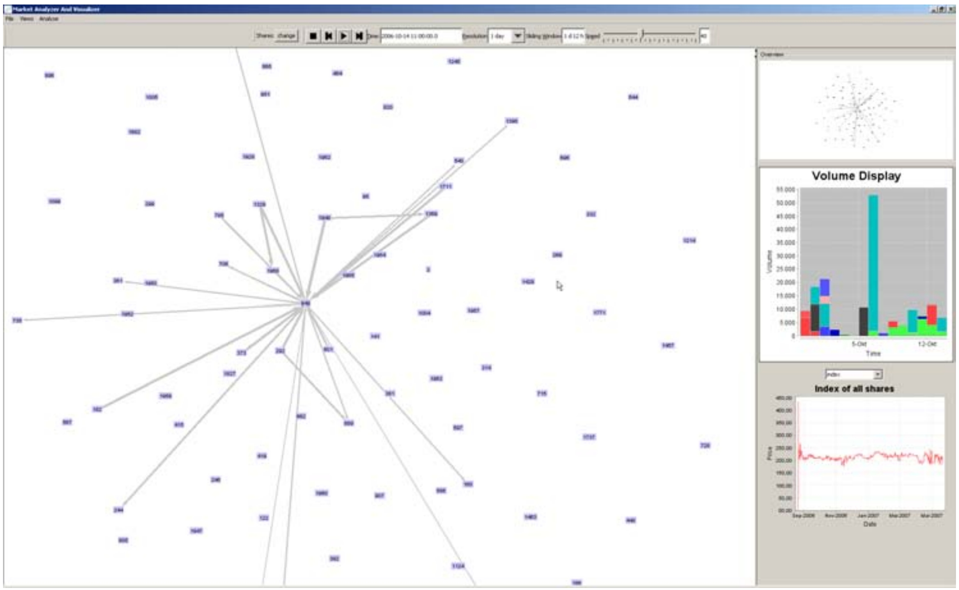

Figure 4.18: Displaying all edges for node 949

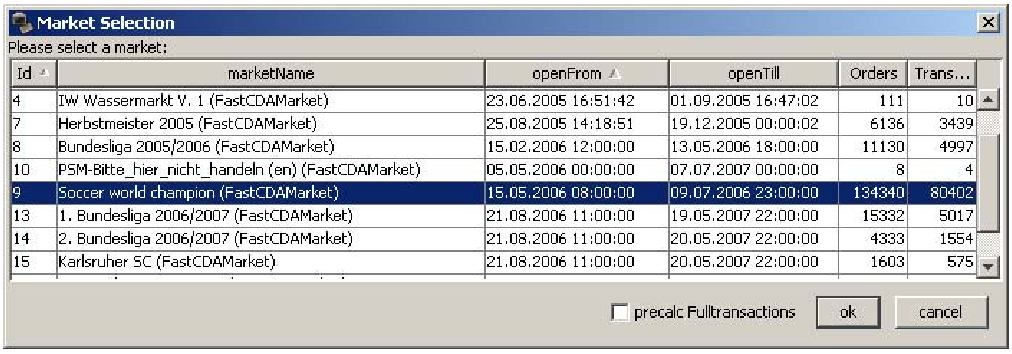

Market Selection

Please select a market:

Id	marketName	openFrom	openTill	Orders	Trans...
4	IW Wassermarkt V. 1 (FastCDAMarket)	23.06.2005 16:51:42	01.09.2005 16:47:02	111	10
7	Herbstmeister 2005 (FastCDAMarket)	25.08.2005 14:18:51	19.12.2005 00:00:02	6136	3439
8	Bundesliga 2005/2006 (FastCDAMarket)	15.02.2006 12:00:00	13.05.2006 18:00:00	11130	4997
10	PSM-Bitte_hier_nicht_handeln (en) (FastCDAMarket)	05.05.2006 00:00:00	07.07.2007 00:00:00	8	4
9	Soccer world champion (FastCDAMarket)	15.05.2006 08:00:00	09.07.2006 23:00:00	134340	80402
13	1. Bundesliga 2006/2007 (FastCDAMarket)	21.08.2006 11:00:00	19.05.2007 22:00:00	15332	5017
14	2. Bundesliga 2006/2007 (FastCDAMarket)	21.08.2006 11:00:00	20.05.2007 22:00:00	4333	1554
15	Karlsruher SC (FastCDAMarket)	21.08.2006 11:00:00	20.05.2007 22:00:00	1603	575

Figure 4.19: Market selection dialog

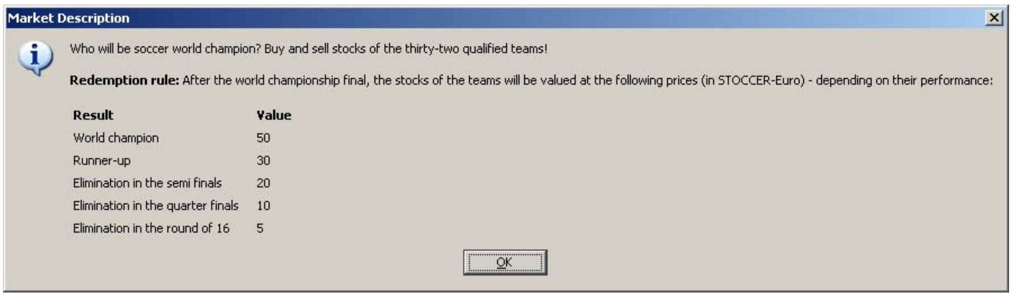

Who will be soccer world champion? Buy and sell stocks of the thirty-two qualified teams!

Redemption rule: After the world championship final, the stocks of the teams will be valued at the following prices (in STOCCER-Euro) - depending on their performance:

Result	Value
World champion	50
Runner-up	30
Elimination in the semi finals	20
Elimination in the quarter finals	10
Elimination in the round of 16	5

Figure 4.20: Market description panel

If a market has never been investigated before, the so-called *fulltransactions* should be pre-calculated. The underlying platform market database is usually normalized to the third normal form. Further information about database normalization can be found in Lang and Lockemann (1995). The reason for the third normal form is the high write performance and reduction of duplicate values. On the other hand, a normalized scheme causes poor reading performance and is far too slow for real time analysis or visualization. To circumvent this, the *precalc* option joins several tables about users and transactions together into a single new table. This significantly increases the ex post analysis performance. The pre-calculation time depends on the database system, server performance, and the number of entries. The precalc function takes about 40 seconds on a Dell Precision M70 laptop using a postgreSQL 8.0 database with about 2000 users, 132,000 transactions, and 246,000 orders. This still seems an acceptable amount of time to start the system.

4.4.2 Market Description

All market participants get certain information about the market. For example, the trading or redemption rules are explained. This information can be very valuable for the investigator to understand certain behavioral differences on different markets (e.g. price jumps according to the different redemption rules). With the prototype the user can open an information panel which shows this information text. Figure 4.20 depicts a screenshot of this dialog.

4.4.3 Account Registration Data

Whenever a new account is created on the platform, the account holder has to enter some information about himself. This is not necessary for trading, but the general terms and conditions state that the user is only allowed to have one single account (c.f. Chapter C). The platform only checks the uniqueness of the username and email address. The most important information (username, forename, surname, email addressee, id, and registration time stamp) is displayed in the account inspector. If the prototype user

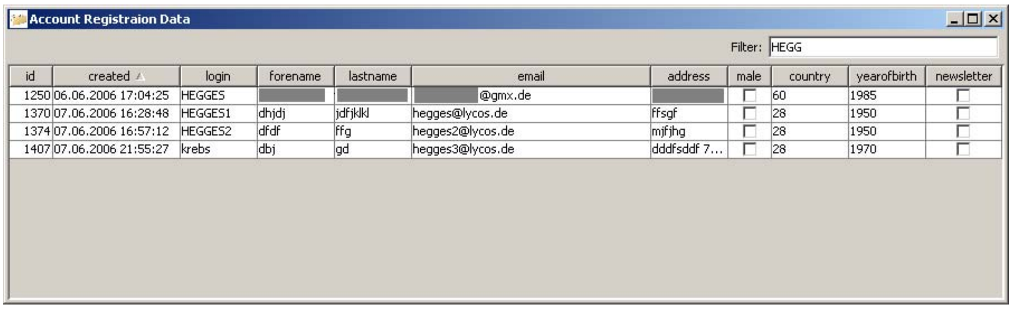

Filter: HEGG

id	created	login	forename	lastname	email	address	male	country	yearofbirth	newsletter
1250	06.06.2006 17:04:25	HEGGES	[redacted]	[redacted]	[redacted]@gmx.de	[redacted]	□	60	1985	□
1370	07.06.2006 16:28:48	HEGGES1	dhjdj	jdfjklkl	hegges@lycos.de	ffsgf	□	28	1950	□
1374	07.06.2006 16:57:12	HEGGES2	dfdf	ffg	hegges2@lycos.de	mjfjhg	□	28	1950	□
1407	07.06.2006 21:55:27	krebs	dbj	gd	hegges3@lycos.de	dddfsddf 7...	□	28	1970	□

Figure 4.21: Account registration data dialog

wanted to check whether someone registered with similar data, he can open the user registration data dialog (Figure 4.21).

In the upper right corner, an integrated search text field can perform an incremental search upon various fields. As shown in Figure 4.21, the search is case insensitive. If the user, for example, types the letters "hegg" all rows are displayed that either have the character sequence in the username, forename, last name, email or address. The above example is taken from the STOCCER data set. The user *HEGGES* registered on June 06th, 2006 with what is probably his real name and address (anonymized in the screenshot). On June 07th, he created two secondary accounts who were manually found and blocked. On the same day at 10 pm, the user created another secondary account *krebs* which had not been found and blocked during the run time of the market. The search list in the prototype reveals the relationship. Note, the user has to type only the four letters "hegg" to verify his suspicion. The user is encouraged to investigate since the username *krebs* is also reported by the algorithms from Chapter 2.

4.4.4 Account Ranking

One of the incentives for the participants in STOCCER is a high score listing on the web page. The ranking is determined by the value of their portfolio. The higher the current portfolio value according to the current market prices, the higher the username appears in the ranking. The ranking is calculated once a day and the result is stored in a database table.

Besides the current portfolio value, the change in the value from the day before is stored in the database. The dialog is similar to the account registration data dialog (see Figure 4.22). Since the screenshot has been taken some weeks after the market had already closed, the *change* column contains only zeros. The *filter* text field allows the user to filter for certain account name patterns. Upon entering a character, an incremental case insensitive search is performed.

The account ranking is important to the investigator because it allows him to gather information about how successful a certain account was. Depending on the account

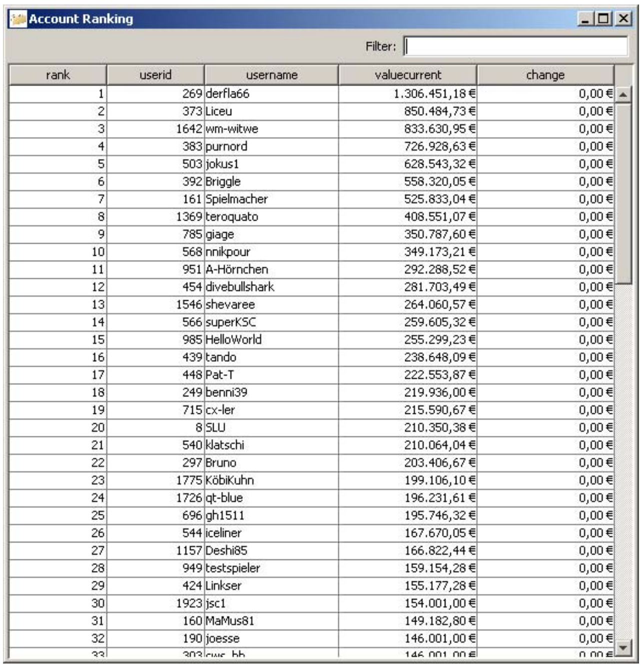

Filter:

rank	userid	username	valuecurrent	change
1	269	derfla66	1.306.451,18 €	0,00 €
2	373	Liceu	850.484,73 €	0,00 €
3	1642	wm-witwe	833.630,95 €	0,00 €
4	383	purnord	726.928,63 €	0,00 €
5	503	jokus1	628.543,32 €	0,00 €
6	392	Briggle	558.320,05 €	0,00 €
7	161	Spielmacher	525.833,04 €	0,00 €
8	1369	teroquato	408.551,07 €	0,00 €
9	785	giage	350.787,60 €	0,00 €
10	568	nnikpour	349.173,21 €	0,00 €
11	951	A-Hörnchen	292.288,52 €	0,00 €
12	454	divebullshark	281.703,49 €	0,00 €
13	1546	shevaree	264.060,57 €	0,00 €
14	566	superKSC	259.605,32 €	0,00 €
15	985	HelloWorld	255.299,23 €	0,00 €
16	439	tando	238.648,09 €	0,00 €
17	448	Pat-T	222.553,87 €	0,00 €
18	249	benni39	219.936,00 €	0,00 €
19	715	cx-ler	215.590,67 €	0,00 €
20	8	SLU	210.350,38 €	0,00 €
21	540	klatschi	210.064,04 €	0,00 €
22	297	Bruno	203.406,67 €	0,00 €
23	1775	KöbiKuhn	199.106,10 €	0,00 €
24	1726	qt-blue	196.231,61 €	0,00 €
25	696	gh1511	195.746,32 €	0,00 €
26	544	iceliner	167.670,05 €	0,00 €
27	1157	Deshi85	166.822,44 €	0,00 €
28	949	testspieler	159.154,28 €	0,00 €
29	424	Linkser	155.177,28 €	0,00 €
30	1923	jsc1	154.001,00 €	0,00 €
31	160	MaMus81	149.182,80 €	0,00 €
32	190	joesse	146.001,00 €	0,00 €
33	303	cwc_bb	146.001,00 €	0,00 €

Figure 4.22: Account ranking dialog

situation, some transactions may even be explained by the expected rise in the ranking. Remember that in STOCCER only the first 100 accounts took part in the final lottery.

4.5 Implementation Design

In this section, an overview of the implementation will be given without presenting the source code or classes. Instead, the focus will be on the general aspects of the implementation.

The prototype is implemented in *Java* using the language features of *Java* 1.5, such as the newly introduced generics, which allow shorter and more elegant implementation of the data handling. Most data is stored in the market platform database. STOCCER used a *PostgreSQL* 8.0 database. *JDBC* is used on the lowest application layer to connect to the market platform. The database is accessed by the implementations of the general market data access facade. The advantage of the facade is that it allows the possibility to adapt the prototype easily to other market platforms. The user has only to implement the facade's interface for the other market platform database scheme. So far, adapters for the *PSM* database scheme and for the *meet2trade* database schema have been implemented.

To reduce the complexity of the UI-classes, the model view controller concept was used for all the graphical components. The prototype uses three open source frameworks to facilitate the necessary extensions of already known and existing components. The *prefuse*[1] framework (Heer and Boyd 2005) is used for the main graph display as well

[1] prefuse `http://prefuse.org`

as the map overview. It provides a Barnes-Hut graph layout implementation. Therefore only the modifications described in chapter 3 had to be implemented.

For the table views, the *glazed lists* open source library[2] from public objects has been used. The library extends the swing default table views and is based on the model-view-controller concept. The *glazed lists* framework was extended in several aspects. The background analysis of the items was introduced for coloring. Furthermore, the mouse and keyboard interaction, formater for time and currencies, and the integration of the *JFreeChart* library had to be introduced.

Finally, the char library *JFreeChart*[3] is used to provide stacked bar charts, pie charts, and scatter plots. All these components are integrated into a swing user interface to assure the portability of the prototype. Debugging is facilitated by extensive logging using the *apache log4j* library.

4.6 Selected Application Examples

In the following, two examples illustrate how the prototype can be used to detect fraud or explain market effects.

4.6.1 Observing Uncommon Trading Volumes

When there is no alert or open case to investigate, the user may just want to observe the current trading activity. Besides the graph of the active traders, the price and volume diagrams allow further insights. If, for example, the trading volume of a certain share rises significantly, the question occurs whether there is new information about the respective share or whether an uninformed trader is investing more than usual.

An example screenshot of such a situation is given in Figure 4.23 on the next page. In the volume display in the lower right, a share is suddenly traded in uncommonly high volumes. Different colors are used to distinguish between different shares. It gets obvious that the significant part of the high volume is based on trades of one single share.

After the user becomes aware of this unusual event, he probably wants to know which share it was and which accounts have been involved. The first question can be answered quickly by moving the mouse cursor over the stacked bar of interest. The tooltip indicates the name of the corresponding share. The latter question can be answered by looking at the graph display. In the center of the display, the two accounts X and Y have traded significantly more than other active nodes in the same moment. If the user left clicks on the edge, the *Account Inspector* will pop up and highlight all transactions between the two neighbors. In the anonymized screenshot in Figure 4.24, the highlighted area is

[2] glazed lists `http://publicobject.com/glazedlists/`

[3] *JFree Chart* `http://jfree.org`

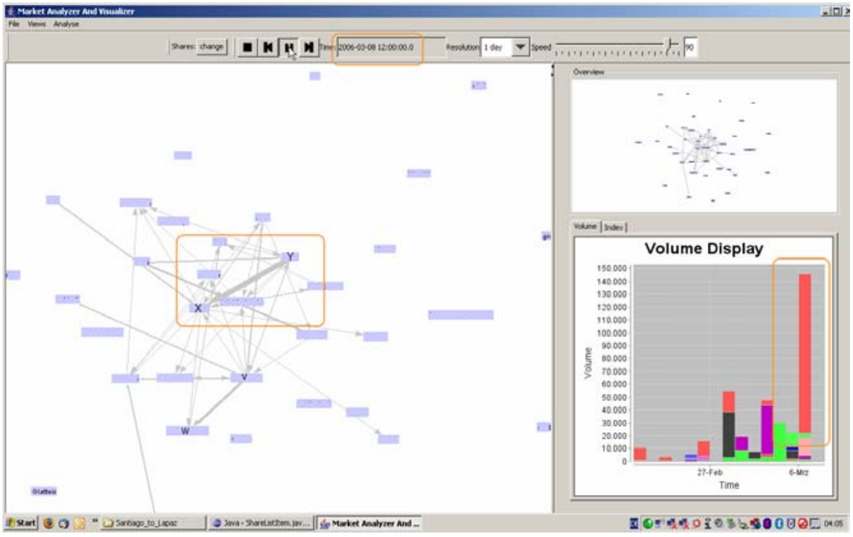

Time: 2006-03-08 12:00:00.0

Resolution: 1 day

Speed: 90

Figure 4.23: Observing uncommon trading volumes

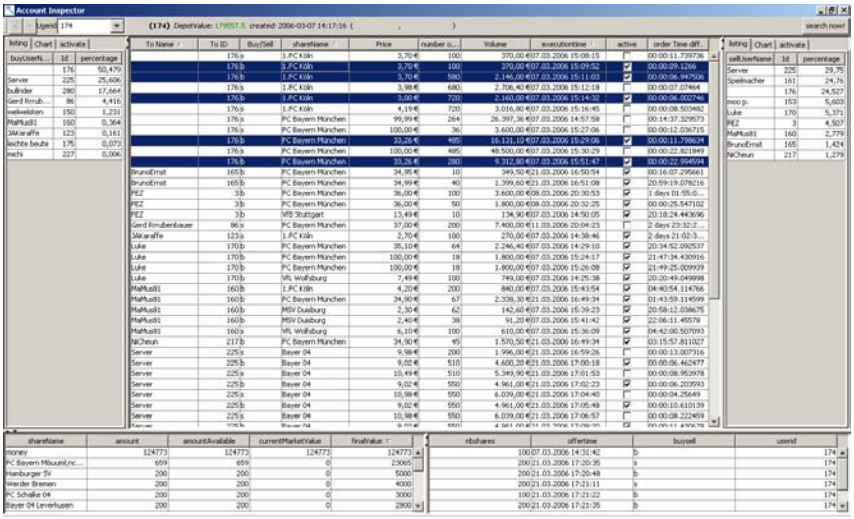

Userid 174

(174) DepotValue: 179557.5 created: 2006-03-07 14:17:16 (,)

buyUserN...	Id	percentage
	176	50,479
Server	225	25,606
bullrider	280	17,664
Gerd Rmub...	86	4,416
welweldken	150	1,231
MaMus81	160	0,364
JAKaraffe	123	0,161
leichte beute	175	0,073
michi	227	0,006

To Name	To ID	Buy/Sell	shareName	Price	number o...	Volume	executiontime	active	order Time diff.
	176	s	1.FC Köln	3,70 €	100	370,00 €	07.03.2006 15:08:15	□	00:00:11.739736
	176	b	1.FC Köln	3,70 €	100	370,00 €	07.03.2006 15:09:52	☑	00:00:09.1266
	176	b	1.FC Köln	3,70 €	580	2.146,00 €	07.03.2006 15:11:03	☑	00:00:06.947506
	176	s	1.FC Köln	3,98 €	680	2.706,40 €	07.03.2006 15:12:18	□	00:00:07.07464
	176	b	1.FC Köln	3,00 €	720	2.160,00 €	07.03.2006 15:14:32	☑	00:00:06.002746
	176	s	1.FC Köln	4,19 €	720	3.016,80 €	07.03.2006 15:16:45	□	00:00:08.503482
	176	s	FC Bayern München	99,99 €	264	26.397,36 €	07.03.2006 14:57:58	□	00:14:37.329573
	176	s	FC Bayern München	100,00 €	36	3.600,00 €	07.03.2006 15:27:06	□	00:00:12.036715
	176	b	FC Bayern München	33,26 €	485	16.131,10 €	07.03.2006 15:29:06	☑	00:00:11.798634
	176	s	FC Bayern München	100,00 €	485	48.500,00 €	07.03.2006 15:30:29	□	00:00:22.821849
	176	b	FC Bayern München	33,26 €	280	9.312,80 €	07.03.2006 15:51:47	☑	00:00:22.994594
BrunoErnst	165	b	FC Bayern München	34,95 €	10	349,50 €	21.03.2006 16:50:54	☑	00:16:07.295661
BrunoErnst	165	b	FC Bayern München	34,99 €	40	1.399,60 €	21.03.2006 16:51:08	☑	20:59:19.078216
PEZ	3	b	FC Bayern München	36,00 €	100	3.600,00 €	08.03.2006 20:30:53	☑	1 days 01:55:0...
PEZ	3	b	FC Bayern München	36,00 €	50	1.800,00 €	08.03.2006 20:32:25	☑	00:00:25.547102
PEZ	3	b	VfB Stuttgart	13,49 €	10	134,90 €	07.03.2006 14:50:05	☑	20:18:24.443696
Gerd Rmubenbauer	86	s	FC Bayern München	37,00 €	200	7.400,00 €	11.03.2006 20:04:23	□	2 days 23:32:2...
JAKaraffe	123	s	1.FC Köln	2,70 €	100	270,00 €	07.03.2006 14:38:46	☑	2 days 21:02:3...
Luke	170	b	FC Bayern München	35,10 €	64	2.246,40 €	07.03.2006 14:29:10	☑	20:34:52.092537
Luke	170	b	FC Bayern München	100,00 €	18	1.800,00 €	07.03.2006 15:24:17	☑	21:47:34.430916
Luke	170	b	FC Bayern München	100,00 €	18	1.800,00 €	07.03.2006 15:26:08	☑	21:49:25.009939
Luke	170	b	VfL Wolfsburg	7,49 €	100	749,00 €	07.03.2006 14:25:38	☑	20:20:49.049898
MaMus81	160	b	1.FC Köln	4,20 €	200	840,00 €	07.03.2006 15:43:54	☑	04:40:54.114766
MaMus81	160	b	FC Bayern München	34,90 €	67	2.338,30 €	21.03.2006 16:49:34	☑	01:43:59.114599
MaMus81	160	b	MSV Duisburg	2,30 €	62	142,60 €	07.03.2006 15:39:23	☑	20:58:12.038675
MaMus81	160	b	MSV Duisburg	2,40 €	38	91,20 €	07.03.2006 15:41:42	☑	22:06:11.45578
MaMus81	160	s	VfL Wolfsburg	6,10 €	100	610,00 €	07.03.2006 15:36:09	☑	04:42:00.507093
NiCheun	217	b	FC Bayern München	34,90 €	45	1.570,50 €	21.03.2006 16:49:34	☑	03:15:57.811027
Server	225	s	Bayer 04	9,98 €	200	1.996,00 €	21.03.2006 16:59:26	□	00:00:13.007316
Server	225	b	Bayer 04	9,02 €	510	4.600,20 €	21.03.2006 17:00:18	☑	00:00:06.462477
Server	225	s	Bayer 04	10,49 €	510	5.349,90 €	21.03.2006 17:01:53	□	00:00:08.953978
Server	225	b	Bayer 04	9,02 €	550	4.961,00 €	21.03.2006 17:02:23	☑	00:00:06.203593
Server	225	s	Bayer 04	10,98 €	550	6.039,00 €	21.03.2006 17:04:40	□	00:00:04.25649
Server	225	b	Bayer 04	9,02 €	550	4.961,00 €	21.03.2006 17:05:48	☑	00:00:10.610139
Server	225	s	Bayer 04	10,98 €	550	6.039,00 €	21.03.2006 17:06:57	□	00:00:08.222459
Server	225	b	Bayer 04	9,02 €	550	4.961,00 €	21.03.2006 17:08:20	☑	00:00:11.430678

sellUserName	Id	percentage
Server	225	29,75
Spielmacher	161	24,76
	176	24,527
moo p.	153	5,603
Luke	170	5,371
PEZ	3	4,507
MaMus81	160	2,779
BrunoErnst	165	1,424
NiCheun	217	1,279

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	124773	124773	124773	124773
FC Bayern M&uuml;nc...	659	659	0	23065
Hamburger SV	200	200	0	5000
Werder Bremen	200	200	0	4000
FC Schalke 04	200	200	0	3000
Bayer 04 Leverkusen	200	200	0	2800

nbshares	offertime	buysell	userid
100	07.03.2006 14:31:42	b	174
200	21.03.2006 17:20:35	s	174
200	21.03.2006 17:20:48	b	174
200	21.03.2006 17:21:11	s	174
100	21.03.2006 17:21:22	b	174
200	21.03.2006 17:21:35	b	174

Figure 4.24: Drilling down onto the relevant transactions

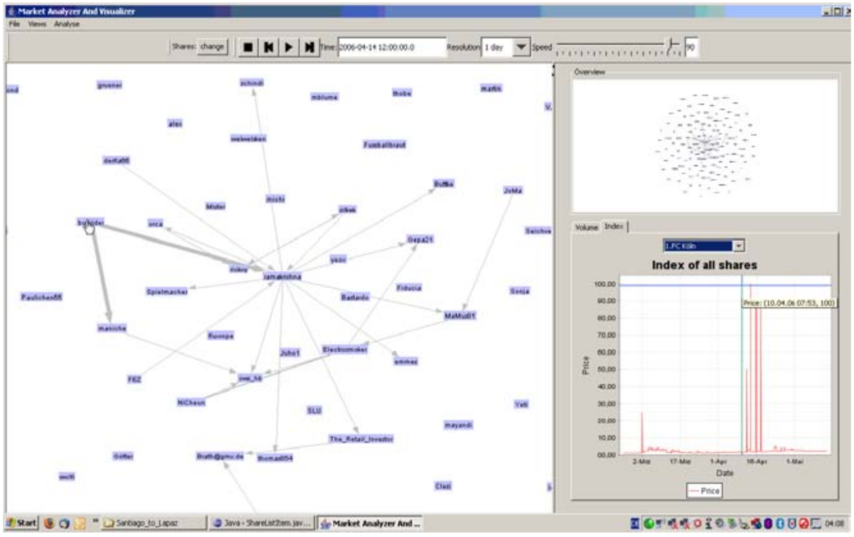

Figure 4.25: Observing price peaks

already reduced to show just the buy transactions. This way the pattern becomes obvious. The pattern is probably related to a feature of the trading screen of STOCCER. The user could click on a position in the orderbook and the order input mask was automatically filled with the corresponding data. So the person probably placed a sell order with the first account inside the spread. Then another session was opened for the second account allowing the open sell order to be bought and a new sell order to be placed on the other side of the spread. These three steps (switch, buy, sell) were repeated again and again. The shortest period between sell and buy orders is 6 seconds. This is very fast for a web-based trading interface and very unlikely to be achieved by accident (besides it is hard to explain why somebody repeatedly played such a ruinous pattern). In the side panels, it is shown that the second account was a very important trading partner for this account. In the table in the lower left, we can see that money and shares of the team *FC Bayern München* are the base of the wealth of this account.

For this investigation, only two actions are necessary: one *mouse over* and one *mouse click*.

4.6.2 Explaining Price Peaks

Besides uncommon trading volumes, price peaks may call the user's attention. The peaks may be observed in the price index or on a selected share. Figure 4.25 shows a price peak example. In the lower right, several peaks up to 100 € can be observed. The time control in the toolbar shows that the graph represents the day of the peaks. In the graph display, two edges appear to be stronger than any other. With a right click on one of these edges, the *Account Inspector* pops up and highlights the transactions between these two accounts (see Figure 4.26).

In this case, the investigator is lucky because the account *maniche* seems to be responsible

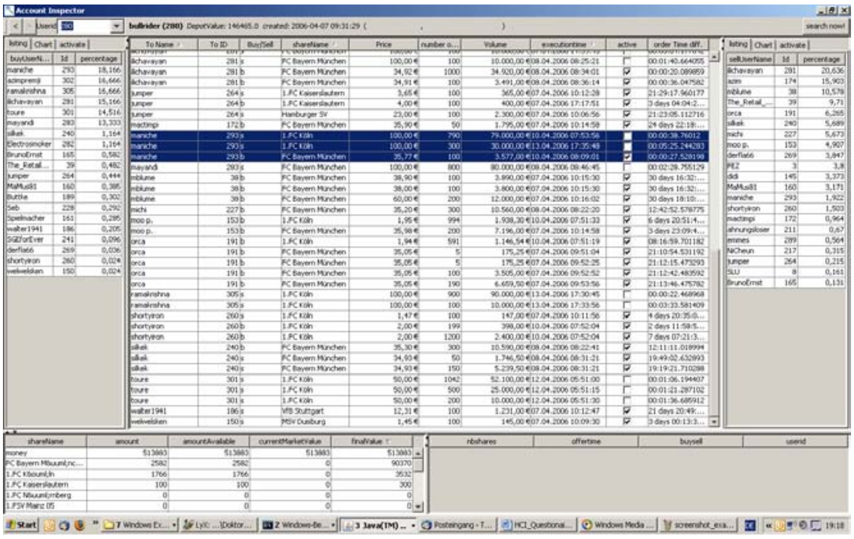

bullrider (280) DepotValue: 146465.0 created: 2006-04-07 09:31:29 (

buyUserN...	Id	percentage
maniche	293	18,166
azimpremji	302	16,666
ramakrishna	305	16,666
ilchavayan	281	15,166
toure	301	14,516
mayandi	283	13,333
silkek	240	1,164
Electrosmoker	282	1,164
BrunoErnst	165	0,582
The_Retail...	39	0,482
jumper	264	0,444
MaMus81	160	0,385
Buttke	189	0,302
Seb	228	0,292
Spielmacher	161	0,285
walter1941	186	0,205
SGEforEver	241	0,096
derfla66	269	0,036
shortyiron	260	0,024
welwelsken	150	0,024

To Name	To ID	Buy/Sell	shareName	Price	number o...	Volume	executiontime	active	order Time diff.
ilchavayan	281	s	FC Bayern München	100,00 €	100	10.000,00 €	08.04.2006 08:25:21	□	00:01:40.664055
ilchavayan	281	b	FC Bayern München	34,92 €	1000	34.920,00 €	08.04.2006 08:34:01	☑	00:00:20.089859
ilchavayan	281	b	FC Bayern München	34,91 €	100	3.491,00 €	08.04.2006 08:36:14	☑	00:00:36.047582
jumper	264	s	1.FC Kaiserslautern	3,65 €	100	365,00 €	07.04.2006 10:12:28	☑	21:29:17.960177
jumper	264	b	1.FC Kaiserslautern	4,00 €	100	400,00 €	07.04.2006 17:17:51	☑	3 days 04:04:2...
jumper	264	s	Hamburger SV	23,00 €	100	2.300,00 €	07.04.2006 10:06:56	☑	21:23:05.112716
mactimpi	172	b	FC Bayern München	35,90 €	50	1.795,00 €	07.04.2006 10:14:58	☑	24 days 22:18:...
maniche	293	s	1.FC Köln	100,00 €	790	79.000,00 €	10.04.2006 07:53:56	□	00:00:38.76012
maniche	293	s	1.FC Köln	100,00 €	300	30.000,00 €	13.04.2006 17:35:48	□	00:05:25.244283
maniche	293	b	FC Bayern München	35,77 €	100	3.577,00 €	10.04.2006 08:09:01	☑	00:00:27.528198
mayandi	283	s	FC Bayern München	100,00 €	800	80.000,00 €	08.04.2006 08:46:45	□	00:02:28.755129
mblume	38	b	FC Bayern München	38,90 €	100	3.890,00 €	07.04.2006 10:15:30	☑	30 days 16:32:...
mblume	38	b	FC Bayern München	38,00 €	100	3.800,00 €	07.04.2006 10:15:30	☑	30 days 16:32:...
mblume	38	b	FC Bayern München	60,00 €	200	12.000,00 €	07.04.2006 10:16:02	☑	30 days 18:10:...
michi	227	b	FC Bayern München	35,20 €	300	10.560,00 €	08.04.2006 08:22:20	☑	12:42:52.578775
moo p.	153	b	1.FC Köln	1,95 €	994	1.938,30 €	10.04.2006 07:51:33	☑	6 days 20:51:4...
moo p.	153	b	FC Bayern München	35,98 €	200	7.196,00 €	07.04.2006 10:14:58	☑	3 days 23:09:4...
orca	191	b	1.FC Köln	1,94 €	591	1.146,54 €	10.04.2006 07:51:19	☑	08:16:59.701182
orca	191	b	FC Bayern München	35,05 €	5	175,25 €	07.04.2006 09:51:04	☑	21:10:54.531192
orca	191	b	FC Bayern München	35,05 €	5	175,25 €	07.04.2006 09:52:25	☑	21:12:15.473293
orca	191	b	FC Bayern München	35,05 €	100	3.505,00 €	07.04.2006 09:52:52	☑	21:12:42.483592
orca	191	b	FC Bayern München	35,05 €	190	6.659,50 €	07.04.2006 09:53:56	☑	21:13:46.475782
ramakrishna	305	s	1.FC Köln	100,00 €	900	90.000,00 €	13.04.2006 17:30:45	□	00:00:22.468968
ramakrishna	305	s	1.FC Köln	100,00 €	100	10.000,00 €	13.04.2006 17:33:56	□	00:03:33.581409
shortyiron	260	s	1.FC Köln	1,47 €	100	147,00 €	07.04.2006 10:11:56	☑	4 days 20:35:0...
shortyiron	260	b	1.FC Köln	2,00 €	199	398,00 €	10.04.2006 07:52:04	☑	2 days 11:58:5...
shortyiron	260	b	1.FC Köln	2,00 €	1200	2.400,00 €	10.04.2006 07:52:04	☑	7 days 07:21:3...
silkek	240	b	FC Bayern München	35,30 €	300	10.590,00 €	08.04.2006 08:22:41	☑	12:11:11.018994
silkek	240	s	FC Bayern München	34,93 €	50	1.746,50 €	08.04.2006 08:31:21	☑	19:49:02.632893
silkek	240	s	FC Bayern München	34,93 €	150	5.239,50 €	08.04.2006 08:31:21	☑	19:19:21.710288
toure	301	s	1.FC Köln	50,00 €	1042	52.100,00 €	12.04.2006 05:51:00	□	00:01:06.194407
toure	301	s	1.FC Köln	50,00 €	500	25.000,00 €	12.04.2006 05:51:15	□	00:01:21.287102
toure	301	s	1.FC Köln	50,00 €	200	10.000,00 €	12.04.2006 05:51:30	□	00:01:36.685912
walter1941	186	s	VfB Stuttgart	12,31 €	100	1.231,00 €	07.04.2006 10:12:47	☑	21 days 20:49:...
welwelsken	150	s	MSV Duisburg	1,45 €	100	145,00 €	07.04.2006 10:09:30	☑	3 days 00:13:3...

sellUserName	Id	percentage
ilchavayan	281	20,636
azim	174	15,903
mblume	38	10,578
The_Retail_...	39	9,71
orca	191	6,265
silkek	240	5,689
michi	227	5,673
moo p.	153	4,907
derfla66	269	3,847
PEZ	3	3,8
didi	145	3,373
MaMus81	160	3,171
maniche	293	1,922
shortyiron	260	1,503
mactimpi	172	0,964
ahnungsloser	211	0,67
emmes	289	0,564
NiCheun	217	0,315
jumper	264	0,215
SLU	8	0,161
BrunoErnst	165	0,131

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	513883	513883	513883	513883
FC Bayern M&uuml;nc...	2582	2582	0	90370
1.FC K&ouml;ln	1766	1766	0	3532
1.FC Kaiserslautern	100	100	0	300
1.FC N&uuml;rnberg	0	0	0	0
1.FSV Mainz 05	0	0	0	0

nbshares	offertime	buysell	userid

Figure 4.26: Drilling down the two transactions

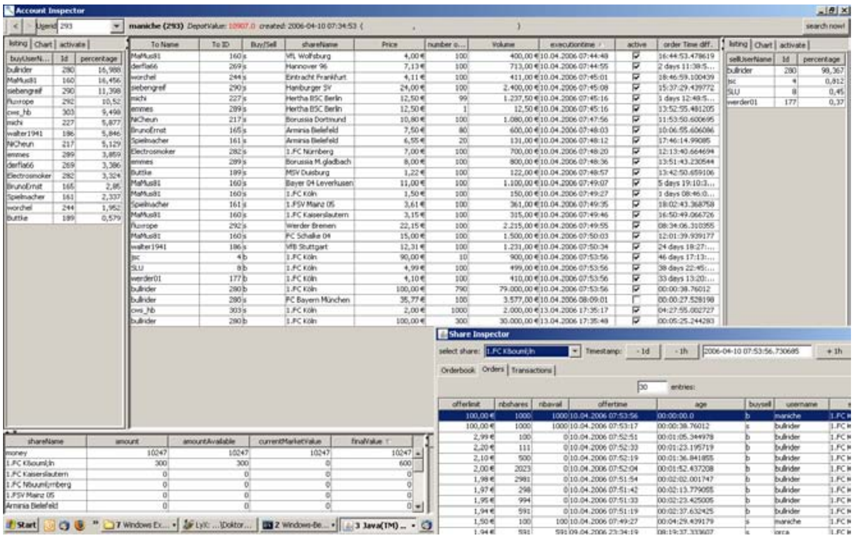

maniche (293) DepotValue: 10907.0 created: 2006-04-10 07:34:53 (

buyUserN...	Id	percentage
bullrider	280	16,988
MaMus81	160	16,456
siebengreif	290	11,398
Fluxrope	292	10,52
cws_hb	303	9,498
michi	227	5,877
walter1941	186	5,846
NiCheun	217	5,129
emmes	289	3,859
derfla66	269	3,386
Electrosmoker	282	3,324
BrunoErnst	165	2,85
Spielmacher	161	2,337
worchel	244	1,952
Buttke	189	0,579

To Name	To ID	Buy/Sell	shareName	Price	number o...	Volume	executiontime	active	order Time diff.
MaMus81	160	s	VfL Wolfsburg	4,00 €	100	400,00 €	10.04.2006 07:44:48	☑	16:44:53.478619
derfla66	269	s	Hannover 96	7,13 €	100	713,00 €	10.04.2006 07:44:55	☑	2 days 11:38:5...
worchel	244	s	Eintracht Frankfurt	4,11 €	100	411,00 €	10.04.2006 07:45:01	☑	18:46:59.100439
siebengreif	290	s	Hamburger SV	24,00 €	100	2.400,00 €	10.04.2006 07:45:08	☑	15:37:29.439772
michi	227	s	Hertha BSC Berlin	12,50 €	99	1.237,50 €	10.04.2006 07:45:16	☑	1 days 12:48:5...
emmes	289	s	Hertha BSC Berlin	12,50 €	1	12,50 €	10.04.2006 07:45:16	☑	13:52:55.481205
NiCheun	217	s	Borussia Dortmund	10,80 €	100	1.080,00 €	10.04.2006 07:47:56	☑	11:53:50.600695
BrunoErnst	165	s	Arminia Bielefeld	7,50 €	80	600,00 €	10.04.2006 07:48:03	☑	10:06:55.606086
Spielmacher	161	s	Arminia Bielefeld	6,55 €	20	131,00 €	10.04.2006 07:48:12	☑	17:46:14.99085
Electrosmoker	282	s	1.FC Nürnberg	7,00 €	100	700,00 €	10.04.2006 07:48:20	☑	12:13:40.664694
emmes	289	s	Borussia M.gladbach	8,00 €	100	800,00 €	10.04.2006 07:48:36	☑	13:51:43.230544
Buttke	189	s	MSV Duisburg	1,22 €	100	122,00 €	10.04.2006 07:48:57	☑	13:42:50.659106
MaMus81	160	s	Bayer 04 Leverkusen	11,00 €	100	1.100,00 €	10.04.2006 07:49:07	☑	5 days 19:10:3...
MaMus81	160	s	1.FC Köln	1,50 €	100	150,00 €	10.04.2006 07:49:27	☑	1 days 08:46:0...
Spielmacher	161	s	1.FSV Mainz 05	3,61 €	100	361,00 €	10.04.2006 07:49:35	☑	18:02:43.368758
MaMus81	160	s	1.FC Kaiserslautern	3,15 €	100	315,00 €	10.04.2006 07:49:46	☑	16:50:49.066726
Fluxrope	292	s	Werder Bremen	22,15 €	100	2.215,00 €	10.04.2006 07:49:55	☑	08:34:06.310355
MaMus81	160	s	FC Schalke 04	15,00 €	100	1.500,00 €	10.04.2006 07:50:03	☑	12:01:39.939177
walter1941	186	s	VfB Stuttgart	12,31 €	100	1.231,00 €	10.04.2006 07:50:34	☑	24 days 18:27:...
jsc	4	b	1.FC Köln	90,00 €	10	900,00 €	10.04.2006 07:53:56	☑	46 days 17:13:...
SLU	8	b	1.FC Köln	4,99 €	100	499,00 €	10.04.2006 07:53:56	☑	38 days 22:45:...
werder01	177	b	1.FC Köln	4,10 €	100	410,00 €	10.04.2006 07:53:56	☑	33 days 13:20:...
bullrider	280	b	1.FC Köln	100,00 €	790	79.000,00 €	10.04.2006 07:53:56	☑	00:00:38.76012
bullrider	280	s	FC Bayern München	35,77 €	100	3.577,00 €	10.04.2006 08:09:01	□	00:00:27.528198
cws_hb	303	s	1.FC Köln	2,00 €	1000	2.000,00 €	13.04.2006 17:35:17	☑	04:27:55.002727
bullrider	280	b	1.FC Köln	100,00 €	300	30.000,00 €	13.04.2006 17:35:48	☑	00:05:25.244283

sellUserName	Id	percentage
bullrider	280	98,367
jsc	4	0,812
SLU	8	0,45
werder01	177	0,37

shareName	amount	amountAvailable	currentMarketValue	finalValue
money	10247	10247	10247	10247
1.FC K&ouml;ln	300	300	0	600
1.FC Kaiserslautern	0	0	0	0
1.FC N&uuml;rnberg	0	0	0	0
1.FSV Mainz 05	0	0	0	0
Arminia Bielefeld	0	0	0	0

select share: 1.FC K&ouml;ln Timestamp: 2006-04-10 07:53:56.730685

30 entries:

offerlimit	nbshares	nbavail	offertime	age	buysell	username	
100,00 €	1000	1000	10.04.2006 07:53:56	00:00:00.0	b	maniche	1.FC K
100,00 €	1000	1000	10.04.2006 07:53:17	00:00:38.76012	s	bullrider	1.FC K
2,99 €	100	0	10.04.2006 07:52:51	00:01:05.344978	b	bullrider	1.FC K
2,20 €	111	0	10.04.2006 07:52:33	00:01:23.195719	b	bullrider	1.FC K
2,10 €	500	0	10.04.2006 07:52:19	00:01:36.841855	b	bullrider	1.FC K
2,00 €	2023	0	10.04.2006 07:52:04	00:01:52.437208	b	bullrider	1.FC K
1,98 €	2981	0	10.04.2006 07:51:54	00:02:02.001747	b	bullrider	1.FC K
1,97 €	298	0	10.04.2006 07:51:42	00:02:13.779055	b	bullrider	1.FC K
1,95 €	994	0	10.04.2006 07:51:33	00:02:23.425005	b	bullrider	1.FC K
1,94 €	591	0	10.04.2006 07:51:19	00:02:37.632425	b	bullrider	1.FC K
1,50 €	100	100	10.04.2006 07:49:27	00:04:29.439179	s	maniche	1.FC K
1,94 €	591	591	09.04.2006 23:34:19	08:19:37.333607	s	orca	1.FC K

Figure 4.27: Switching to *maniche*

for two peaks. But the crucial role seems to be played by *bullrider*, who repeated that pattern again with *ramakrishna* (c.f. transaction list in Figure 4.26). It is highly suspicious that someone would be willing to buy a share for 100 € which had a common market value lower than 3 € and that they would accept this offer within seconds. In total, this account bought *FC Köln* for 109,000.00 €. Furthermore, the offer for 100 € was only seconds old (displayed in the column *order Time diff.*) when *maniche* placed the order. With a double click on *maniche* the *Account Inspector* switches to a view of that account (see Figure 4.27). With this information, now the suspicion turns into certainty. The account *maniche* was created on the same day, sold all other shares, and bought with one order of 1,000 shares at 100 € the rest of the spread and as much as he could of the offer from *bullrider*. The hypothesis that he might be just a big supporter of the share can be checked by opening the *Share Inspector* via the context menu. The *Share*

Inspector shown in the lower right in Figure 4.27 reveals that *bullrider* opened the spread as much as he could just before the order from *maniche*. This way, he gathered enough shares to sell way over price to *maniche*.

For this investigation, only three interactions were necessary: one *mouse click*, a *double click*, and one *context menu* interaction.

In case of a peak in the index chart, the relevant share has to be determined first. Therefore, the user can reveal the last orders in the moment of the peak using the *open order history* option from the context menu. After that, the investigation continues as described above.

4.7 Summary

In this chapter, the prototype has been presented with its different views and features. Besides the general views for each market aspect, the interaction and special dialogs have been described. Finally, a quick implementation overview has been given.

Although one might get the same information with a solid knowledge of SQL, a database query interface, and paper and pencil, it is much more convenient to get the information presented in the way in which it is needed. This is illustrated by the two examples at the end of the chapter, where the user could get the necessary information to make a decision within a few clicks. Although the examples seem to be constructed, they were taken from real world data sets. This underlines the relevance and necessity of such a tool for market operations.

5 Evaluation

In this chapter, the proposed algorithms and the prototype will be evaluated. The visualization has already been addressed in chapter 3.3 on page 64, but it is indirectly touched by the evaluation of the prototype.

Since the main question is how to detect fraud in online auctions, the evaluation of the algorithms is the primary concern. Classifying accounts into fraudsters and normal users is difficult and subjective. While there are clear market rules which the system in most cases does not allow to viloate, there are also possible transactions which are ex post not an optimal decision. It is now up to the judgement of the market operator to decide whether this was intentional or accidental. Since the true classification of the accounts into fraudsters and normal users using just transactional data is such a hard problem, the best approach is to get a self-classification of the account holders themselves. Therefore, a post field study online survey was done.

Furthermore, a human-computer-interface (HCI) experiment was used to evaluate the prototype as well as the algorithms. Certainly, the algorithm evaluation is not the primary focus of an HCI experiment. To include not only the prototype evaluation but also the algorithmic verification in the experiment, the participants had to solve tasks that were designed to objectify the results of the algorithms.

The chapter is organized as follows: The questionnaire and its results are presented and discussed, followed by the introduction of the experiment and the algorithm evaluation, and finally the evaluation of the visualization is presented.

5.1 Post-Experiment Questionnaire

The direct approach to verify whether or not somebody has violated the terms and conditions is to ask him. If he admits to being a fraudster, it is very likely that he tried to trick the market platform. Note, that the opposite does not necessarily hold true. A person might not openly admit to be a violator of the rules though being a fraudster. Since the terms and conditions stated that if fraud was discovered the prizes have to be returned, all persons who received a prize are less likely to admit anything. Still, several conclusions can be drawn from the results obtained.

5.1.1 Design of the Questionnaire

The questionnaire was the second post experiment questionnaire to which the users have been invited. The first one, with a different topic, was conducted immediately after the field study finished in July 2006. The second online survey started in May 2007. The questionaire was implemented as a set of web pages storing the results into a database[1]. The url to the questionnaire was included in the invitation text that was sent via email to the participantes. Each URL included a unique identifier in order to track the different participating accounts.

According to Cook et al. (2000) the response rate is influenced by several criteria such as incentivation and university sponsoring. Deutskens et al. (2004) found that the response rate depends on follow-ups, incentives, length, and presentation of the questionnaire. According to their results, a short questionnaire incentivised by a lottery of small prices but with a high chance of winning and having one follow up should have the best response rate. Therefore, the online questionnaire was shortened to only one page excluding the "welcome" and "good bye" screen. Since the last contact with the participants was almost nine months before the invitation, the incentive to participate played an important role. Three prizes were raffled among the participants who completed the online survey form. The prizes were vouchers from a well known online book and music store with values of € 100.00, € 75.00 and € 50.00. The probability of winning a prize was high compared to the probability of winning in the STOCCER platform itself, since only a few people would answer a survey nine month after the experiment. However, this was not mentioned in the invitation nor in the questionnaire itself.

To prevent people from using similar fraud strategies as employed on STOCCER, each invitation link to the survey contained an identification code. This code was unique for each invitation. This technique allowed for the detection of users who participated several times for the same account. On the downside, using the technique could not ensure that if a person had illegally created more than one account on STOCCER, he or she could participate one time for each account.

To distract the participants from the real purpose of the inquiry, the questions concerned the trading strategy employed during the championships. Besides the strategy, participants were asked about their username, any possible secondary user names, and the email address.

The questionnaire appearance and technology were chosen to be as simple as possible. According to Dillman (2000), different appearance from respondent to respondent may influence the results or –according to Smith and Leigh (1997)– even prevent participants from reaching the survey (i.e. technology or accessibility barriers). The full version of the questionnaire is included in the appendix (c.f. Appendix B). Participants from other

[1] The questionnaire technology was implemented as an online questionnaire by Jan Schröder from the STOCCER team.

countries were in the minority on the STOCCER platform, so only a German version of the questionnaire was designed and therefore sent only to the German speaking members (the majority of the registered users).

An email, with an invitation to participate that included a link to the questionnaire, was sent to 1602 registered accounts of the trading platform. This subset of the participants was selected from the set of all accounts who had registered for the world championship markets by the following criteria:

- the email address had one of the following suffixes ".de", ".ch" or ".at"
- the country of origin, which the participant had selected during his registration, was either Germany, Austria or Switzerland.

If at least one of the criteria matched, the invitation to the questionnaire was sent to the corresponding account.

5.1.2 Questionnaire Results

Two hundered and fourteen participants completed the online questionnaire. This is a response rate of 13.36 percent. The average response rate for online surveys is usually 20-26 percent according to Kaplowitz et al. (2004) .

The questionnaire asked participants to reveal their *strategy*. In a strict sense of experimental economics, a strategy is a specific action triggered by an external stimulation or situation, such as '*buy stocks of Italy if the price drops below 15 VCU*'. Since the strategy space in this sense was too large to be compiled into a questionnaire, the participants were asked about their trading approach or the source of knowledge that they relied on. However, following the terminology of the questionnaire, these approaches will be called strategy in the following. The below answers were offered (translated from German to English):

1. I closely followed the news and tried to use this information to my advantage.
2. I tried to use my knowledge about soccer and my experience in buying and selling stocks to my advantage
3. I invested all my money into one team from the very beginning.
   a) I created multiple accounts in order to bet different teams.
4. I started by betting on the champion(s) from previous year(s).
5. I invested in outsiders.
6. I frequently traded shares of teams already excluded.
7. I created more than one account in order to concentrate shares and/or money on one account by:

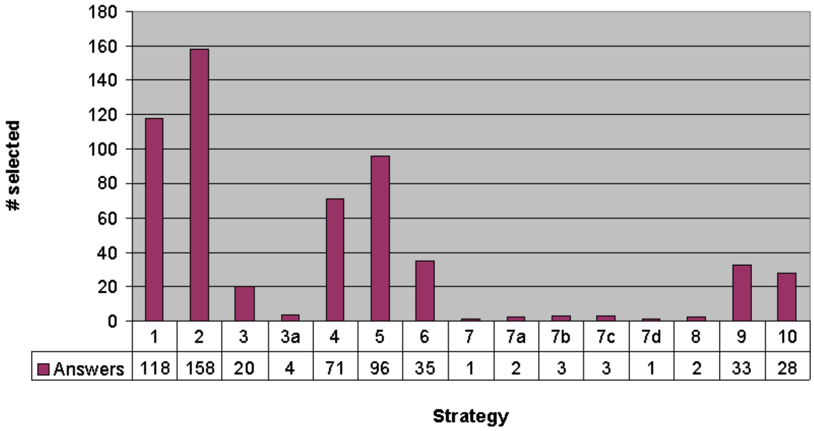

The strategies corresponding to the numbers above can be found in the questionnaire in Appendix B.

Figure 5.1: Absolute number of answers per question in the online questionnaire

a) trading shares back and forth between the two accounts

b) selling overpriced stocks to other accounts.

c) buying underpriced stocks from the other accounts.

d) opening the spread with a third account to use the widened spread for further transactions between two of my accounts.

8. I used a script/bot to trade faster and with less effort

9. I used Excel or another program to keep a better overview about my depot.

10. I used a different strategy (please describe below)

The distribution of the answers is depicted in figure 5.1. Multiple answers were allowed. The most common strategies were number 2 (rely on personal knowledge and experience), 1 (news based trading), 5 (trading outsider teams) and 4 (trading the top teams of the last championships). These strategies are *simple*, *intuitive*, *obvious* and can probably be generalized to many fields of trading on stock exchanges. They are *simple* because they do not rely on any special technology or market mechanisms and assume price changes because of news. They are *intuitive* since they rely on personal knowledge and experience instead of mathematical models and *obvious* because outsider teams are likely cheaper initially but promise the highest returns in case of success.

Less common and more interesting for this evaluation are the people admitting to have used strategy 3a and 6, 7, 7a -7d, 8, and 9. Strategy 6 (trading teams who lost and were out of the champion ships) is a strategy particularly relevant for the tournament market on the STOCCER platform. Since all shares had to add up to a sum of 200 VCU (virtual currency units) and the overall stock exchange should be a zero-sum game, it is difficult to remove a share out of the portfolio without imbalancing the market. While this is still easy for a share that has a final payoff of zero, the later a team dropped off the higher its price was. So to remove the teams who dropped off at the quarter finals would have been quite difficult while keeping the option of portfolio trades.

The high number of results for this strategy is in line with our observation of the methods used on these markets to make further profits. The shares were cheap in absolute terms compared to the teams continuing in the tournament and got less attention from the other traders. Though it is not fraud to trade these shares, buying these shares for a price > 0 indicates either lack of understanding of the rules or having a malicious intention. Certainly, there is a valid reason to sell the shares in case their market value is still higher than their final value. This should drive the market price quickly to the final value. On the STOCCER market platform, the contrary could be observed. Prices of excluded teams could raise a few days after their exclusion again. This could be either with the intention to transfer money to another account by being willing to buy for an unreasonable price or by making sure to be able to sell them again for a even higher price or later on. Note, that this strategey could not only be employed on teams with a definite payout of 0, but any team who had dropped off the competition and therfore had a definitive final payout value. However, the least attention and thus the highest potential gains could be achieved in the zearo payout share markets.

Strategy 9 (used a tool such as MS Excel to get a better overview) was surprisingly high (15.4%). It is an indicator that the design of the trading interface could be further improved. This strategy was not illegal according to the general terms and conditions. But it is highly related to strategy number 8 which referred to the use of a script or robot for trading. Two people admitted having employed scripts or robots. Both accounts had been found to be fraudsters before the survey started using the prototype for manual investigations. Finally, 3a and 7a-d (all concerning the creation of multiple accounts) were mentioned 14 times in total.

5.1.3 Discussion

With 13.36 percent, the response rate is significantly lower than the expected rate of 20-26 percent. Several reasons come into mind about why the response rate was considerably lower than the expected average:

1. The participants might have used a secondary or even temporary mail account such as mailinator, a common technique to prevent spam to one's main mail account, to register for STOCCER. Since they would not expect emails from the platform nine month after the championship finished, they would not be checking this account on a regular basis.

2. Some participants might not be willing to reveal their trading strategy. Especially successful traders might want to keep their strategy secret for future applications.

3. If a person had violated the general terms and conditions by creating several accounts, he might only answer for one of them. Again for reasons similar to the already mentioned under (1) or because he feared that he might get caught and prosecuted (e.g. had to return a prize).
4. After nine month the participants did not care anymore about STOCCER. As described in Chapter 3, many participants only traded once or twice and did not return afterwards.
5. Even if they still had interest in the topic, they might have thought that they could not contribute anything since they did not pursue any concrete trading strategy.

These are all only speculations about the low response rate. But why is the response rate so important? Depending on the explanations above, the real number of participants of the survey may differ. Consequently, and even more relevant, the number of fraudsters has to be estimated significantly higher assuming that fraudsters might have used their multiple identities again with random answers and by this lowering the overall number of participants furthermore. Unfortunately it is almost impossible to verify these different possibilities. But the response rate is still important to approximate the number of fraudsters on STOCCER.

Twelve of the 214 participants voluntarily admitted to being a fraudster (5.6%). Note that this number is lower than the sum of positive answers to the fraud strategies 3a, 7a-d and 8, because one person could select multiple strategies. Furthermore, three accounts that were excluded during the championship because of fraud, did not admit it in the questionnaire. This increases the number of fraudsters to 15 (7.0%).

One of the accounts, who did not admit any fraud, even tried to increase his possibility of winning the survey lottery. He took part in the survey with two of his accounts, not admitting any fraud; for one of the two accounts he even took part twice with different answers to the questions. The danger of fraud in online surveys is already addressed in Smith et al. (1997). They suggest to circumvent the problem by using one time passwords. In the survey presented here, an invitation code has been used as a one time password. Interestingly this user had three different invitation codes, filed the survey on three different days, but used the same username twice. Either he had forgotten for which user he had used the invitation code or he did not care too much to cover his activity. Thus, he must have had at least three accounts. This again confirms two hypothesis:

1. The number of individuals behind the registered accounts who participated in STOCCER and the survey is lower than the absolute number of answers on the questionnaire. So the dark figure of fraudsters is probably even higher than 7%.
2. Online lotteries are an easy target of fraud, since the effort to commit fraud is much lower than in comparable real world settings. In case of an online questionnaire, it is easier to reply multiple times than the comparable actions of completing several

questionnaires, putting them in envelopes, going to the post office, getting stamps and mailing them back to the origin.

Even more interesting is the fact that twenty other persons answered the question regarding which strategy they would use if STOCCER were repeated by marking at least one of the fraud strategies. This is almost twice as much as the initial set of persons admitting fraud. Of the formerly honest market participants, 9.9% would switch to the group of fraudsters. This confirms the bad influence of fraudsters on honest traders if the honest traders become aware of fraudsters on the platform. This underlines the importance of market supervising authorities.

5.2 Human-Computer-Interface Experiment

The accounts classified as fraudsters by the presented detection algorithms are not necessarily all fraudsters, a manual classification was necessary. To objectify this manual classification at least three different persons should give their opinion about each account. A lab experiment with the prototype was done. The experiment was structured as follows:

1. Welcome participant and explain the following steps:
2. Answering one question about the expectations for such a market monitoring and fraud detection system.
3. Watching fifteen minutes of tutorial videos explaining how to use the prototype. Afterwards the participant can ask questions.
4. Solving three tasks related to three different scenarios/situations:
   a) Check thirteen account pairs and tell whether they belong to the same person or have committed fraud. Possible classifications were fraudster, innocent, unsure.
   b) Check a single account to tell whether he has committed fraud. Possible classifications were fraudster, innocent, unsure.
   c) Monitor five days of market activity and list all events that appear strange / not normal.
5. Completing a post-study questionnaire that includes: asking how well the test person's expectation were met by the prototype, rating the perceived task difficulty, and taking a human-computer-interface questionnaire.

The full questionnaire is included In Appendix D. The post-study questionnaire was adopted from Lewis (1995). The initial question asked after the welcome was to compile a user requirement list and compare it to the list compiled from related work (see Chapter

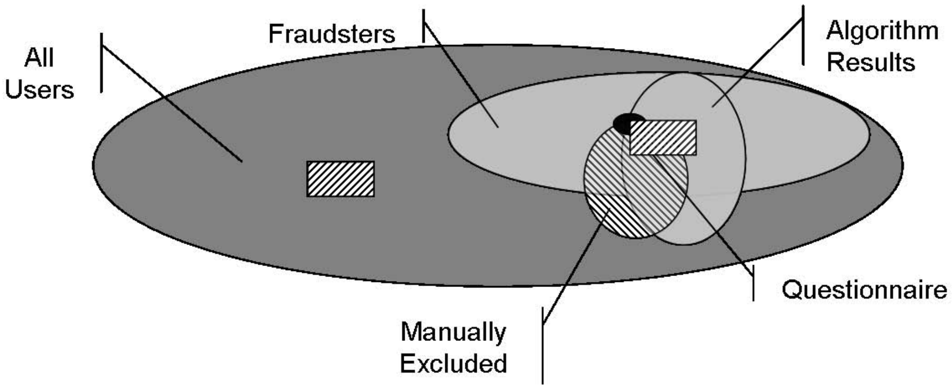

Figure 5.2: Overview of the different account sets

3.2.1). The three tasks were designed to test three different scenarios and started from the easiest scenario, the investigation of a very specific suspicion involving just two accounts, raising to the most complex scenario, where all activity in the market has to be monitored. The amount of information needed for each task was very different. The more information involved the better the tool has to help support the user to make a decission.

The experimenter was familiar with all details of the system and wrote a protocol about all questions, comments, and time the participant needed to complete each task. All participants were familiar with markets and the underlying trading platform STOCCER. The majority were PhD students from the Karlsruhe Institute of Technology. Since only one desktop system was available with a large screen (1920x1200, 24"), all participants completed the experiment one after another. A secondary monitor (1600x1200, 21") expanded the virtual desktop to 2420x1200. Interestingly, it was only used by one participant.

5.2.1 Selection of the Test Data Set

To validate the classification quality of the algorithm in tasks one and two, accounts from four different subsets were selected. Figure 5.2 depicts the subset space. The outer ellipsoid symbolizes the set of all registered accounts. Based on the observation and exclusions during the field study as well as the results from the follow up questionnaire, fraudsters is suspected within this group. Intersecting with the fraudsters and normal users, the group of accounts reported by our algorithms is located. Some of the fraudsters were found and excluded during the run-time of the field study ($\{Manually\, Excluded\}$). All members of this group are supposed to be fraudsters, but the only argument is that they did not complain about being excluded from the market. To understand why this might not have been the case, it is important to remember that many people only logged in once or twice and never came back again. Thus, it is uncertain whether or not the members of this subset even noticed or cared that their accounts had been blocked.

Still, members of this subset have a higher probability of beeing fraudsters than the ones identified by the algorithms. Finally, there is the small group of accounts who admitted fraud in the online survey described above. This is classified with the highest reliability and are represented by the small dark ellipsoid within the set of fraudsters intersecting with the $\{Questionnaire\}$ and $\{Manually\,Excluded\}$ sets.

To make the manual classification a real challenge, accounts from all subsets were selected. This includes accounts that are supposed to be innocent. Ideally, this is the set of all users excluding the fraudsters. Since the fraudster subset is unknown, the best approximation is the set of all users excluding $\{Manually\,Excluded\}$, $\{Algorithm\,Results\}$ and $\{Questionnaire\}$. Since the $\{Questionnaire\}$ subset was significantly smaller, a lower proportion was included in each task.

For the first task, four pairs (eight accounts) were chosen randomly without repetition from each of of the following sets:

- $\{Manually\,Excluded\}$,
- $\{Algorithm\,Results\}\backslash\{Manually\,Excluded\}$ (to prevent double checks)
- $\{Innocent\,Users\}$.

From the $\{Questionnaire\}$ subset only one pair was included because of the small overall size of this subset. Three sets times four pairs plus one pair for the $\{Questionnaire\}$ subset adds up to 13 account pairs.

For the second task, three accounts were chosen randomly from each set, except for the $\{Questionnaire\}$ where only one account was randomly chosen. Thus each experiment participant checked 11 accounts (four pairs plus three individual accounts) from each of the first three subsets and three further accounts from the $\{Questionnaire\}$ subset. The order of the accounts was shuffled randomly for each participant of a group to forestall learning effects.

Overall four groups participated in the experiment. Since each group verified three $\{Questionnaire\}$ accounts and 11 accounts of the other 3 subsets, overall 44 accounts of each subset were tested except for the $\{Questionnaire\}$ set where only 12 accounts were available. Since not all subsets contained 44 accounts, some of the accounts had to be repeated. These were not drawn randomly but explicitly chosen. While in the $\{Manually\,Excluded\}$-set some accounts have been repeated, the $\{Algorithm\,Results\}$ have been extended with some accounts only reported via the Gini algorithm.

5.2.2 Experimental Results

Since the goal of the experiment was twofold, this section is split into two subsections. The first one addresses the results regarding the algorithm precision evaluation while the second part focuses on the prototype usability.

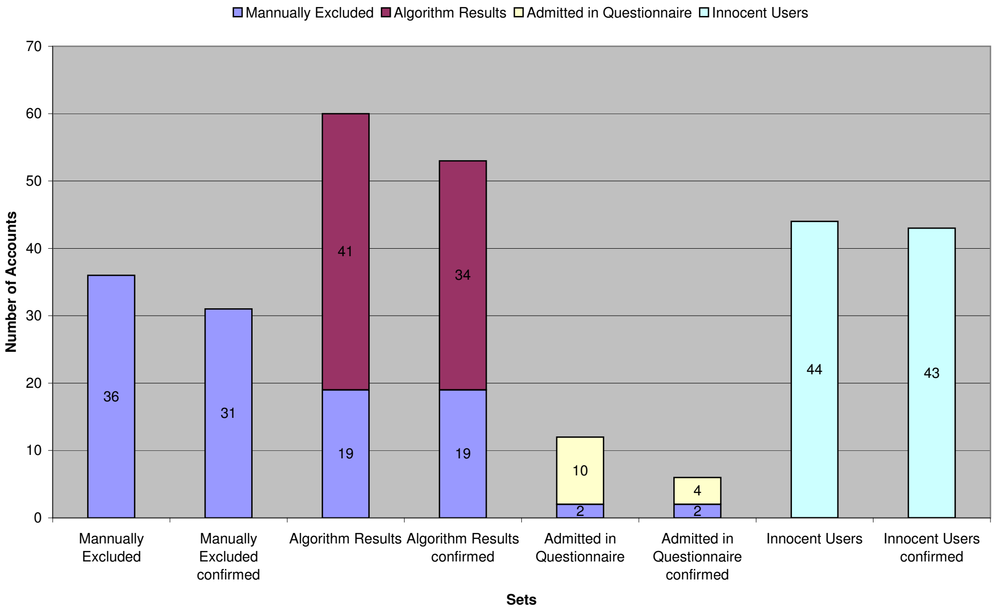

Figure 5.3: Number of accounts confirmed in each set by the experiment

Algorithm Precision

Does the classification of the experiment participant match with the expected one? To answer this question, the three options the participants could use (*honest*, *fraudster*, *unsure*) have to be mapped to the simple classification of *fraudster* and *honest*. The rule employed in this analysis was that an account was only considered to be a *fraudster* if at least two participants had marked him as such.

Figure 5.3 shows as a bar-chart how well the participants agreed with the pre-classification. The first bar of each set indicates the number of accounts used in the experiment of that specific set. For example, the set of $\{InnocentUsers\}$ was much larger, but only 44 of the 1195 accounts were used for the experiement. The second bar with the additional tag *confirmed* indicates how many of the accounts have been classified the same way by the participants (according to the rule given above). The classification was in most cases unambiguous. The ambiguous case of three different opinions (*fraud*, *honest*, *unsure*) or high uncertainty (2 or 3 times *unsure*) appeared twice for the $\{Algorithm\,Results\}$-set, twice for the $\{Innocent\,Users\}$-set and once for $\{Manually\,Excluded\}$-set. They were classified according to the proverb "in dubio pro reo" as innocent users. The interpretation of the low rate of ambiguous cases cannot be generalized, since the sample size is rather small. Still, it indicates how difficult it is to differentiate in some cases between an innocent user and a fraudster even for a human being. One account of the innocent user set got classified by two persons as fraudster (c.f. 5.3). Though the transaction history of the innocent accounts had been previously checked, the registration data

Table 5.1: Fleiss' Kappa and its interpretation

Group	Kappa*	Kappa$^+$
1	0.49	0.53
2	0.31	0.47
3	0.50	0.58
4	0.51	0.61

Kappa	Interpretation
< 0	poor agreement
$0.0 - 0.20$	slight agreement
$0.21 - 0.40$	fair agreement
$0.41 - 0.60$	moderate agreement
$0.61 - 0.80$	substantial agreement
$0.81 - 1.00$	almost perfect agreement

*taking only *normal* into account
$^+$taking *unsure* as *normal* as well

had been omitted due to missing automatism. When revising the registration data with the prototype, the participants found another account with highly similar registration data and the identified account made at least one suspicious transaction at the end of the championship to a third account. This convinced two participants to mark him as fraudster.

The difficulty in revealing a fraudster can also be seen in the poor classification rate of the questionnaire cases. All cases except the one from the $\{Manually\,Excluded\}$-set had no common transactions. They were only related by similar registration data. The participants of the experiment had difficulties finding this link. This indicates room for future improvements.

Regarding the detection precision, the algorithms detected 60 accounts, of which 53 have been confirmed. Thus, the rate is 88% correctly identified or 12% false positives. The manually classified accounts of the $\{Manually\,Excluded\}$-set had a similar, but slightly lower accuracy (31/36=86%). The false negatives cannot be determined exactly. A rough estimation can be done by summing up undetected accounts. In the $\{Manually\,Excluded\}$-set, 12 approved accounts have not been detected; a further 10 in the questionnaire and one in the honest set. This sums up to 23 accounts which is about two percent false negatives (23/1224=1.88%). Since not all accounts have been manually classified this number has to be taken as a lower boundary.

Another aspect is the cosistency between the participants regarding their classifications which is also known as inter-rater reliability. A standard measurement is the Fleiss' Kappa (1971) which is defined in the following way:

$$K = \frac{p_0 - p_e}{1 - p_e} \tag{5.1}$$

with $p_0 - p_e$ being the degree of agreement actually achieved and $1 - p_e$ the degree of agreement attainable by chance. The standard interpretation of K is shown in Table 5.1. The values for the groups varied between 0.31 and 0.51 which is between fair and moderate agreement. The value is slightly lower since the users had the option to classify an account as unsure. If –in dubio pro reo– the unsure classifications are interpreted as

Table 5.2: Size, number of transactions, and volumes per group

	all accounts	survey participants	experiment accounts	all discovered fraudsters	excluded fraudsters
group size	1260	187	127	62	36
# transactions	80402	23038	25739	15715	1761
volume (m VCU)	225.4	66.6	86.3	56.8	12.2

normal users, then the agreement improves to between moderate and substantial (c.f. Table 5.1, column *unsure as normal*).

Market Impact of the Fraudsters

Another interesting aspect are the shares of the different groups addressed in the experiment. Table 5.2 shows the group size as number of accounts, number of transactions and the volume (shares*price) of the different sets. The whole market is given as a reference. Almost a third in terms of volume participated in the strategy survey and more than a third of the transaction volume was covered by the experiment. The second to last column refers to all discovered fraudsters (by hand and algorithm) while the last column refers to all fraudsters excluded during market operations. This means that 25% of the overall transaction volume belonged to fraudsters but only 5.4% have been detected manually. The algorithms discovered almost five times more transactions (11865) and almost 4 times more volume (43.9m VCU) than the manual classifications.

5.3 Visualization

The evaluation of a prototype is a difficult task. Usually there exists neither a similar system to compare to, nor have the testers enough experience to think “out of the box” for substantial criticism. In the field of visualization it is even harder, since visual disabilities (e.g. color blindness) of potential users have to be taken into consideration. Furthermore people have different tastes about colors and appearances in general. So personalization is important, too.

The main focus of software evaluation is the requirement compliance, usability and number of bugs. During the experiment only a small number of bugs occurred – most of them inconsistent labeling. Therefore, the number of bugs will not be further addressed. The term usability shall be understood in the following way:

Definition. Usability

Usability is defined as the extent to which a product can be used by specified users to achieve specified goals with effectiveness, efficiency and satisfaction in a specified context of use (ISO 9241, 1996).

In this chapter, the requirement compliance and usability of the prototype shall be addressed. To test for compliance, especially focusing on the visualization, different visualization taxonomies will be discussed first.

5.3.1 Information Visualization Taxonomies

Early work on visualization taxonomies date back to 1990. Roth and Mattis (1990) developed a catalog of objects and tasks together with suitable visualizations. They developed a data analysis and visualization tool called SAGE, which is an automatic presentation system. In a followup, Wehrend and Lewis (1990) refined the taxonomy of Roth and Mattis by changing the operations (c.f. table 5.4). Object classes (c.f. table 5.3) and operation classes are the two axis of the taxonomy.

Table 5.3: Object classes (Roth and Mattis 1990)

scalar
scalar field
nominal
direction
direction field
shape
position
spatially extended region or object
structure

Table 5.4: Operation classes (Wehrend 1990)

identify
locate
distinguish
categorize
cluster
distribution
rank
compare
within and between relations
associate
correlate

Every complex task has to be broken down into simple objects and the related task into basic operations fitting the taxonomy. Then appropriate visualizations can be derived from the matrix of Roth and Mattis.

Though a designer can verify whether the chosen visualization fits the operation and data, user interface design includes more than just the pure visualization. Based on Nielson et al. (1990) nine basic usability principles should be tested. Though these principals seem obvious they are often hard to apply in practice as Nielsen and Molich write (1990).

1. Simple and natural language
2. Speak the user's language
3. Minimize user memory load
4. Be consistent
5. Provide feedback
6. Provide clearly marked exits
7. Provide shortcuts
8. Good error messages
9. Prevent errors

Though these principles may serve as guidelines for application design in general, small differences can have a significant influence on user perception and performance. Kobsa (2001) made an empirical comparison of three commercial information visualization systems. This comparison showed that even though the systems offered similar views, small features like the default visualization or ease of use of the possible interactions influenced the results significantly. The users of one system solved the given tasks faster, while users of a different system achieved higher accuracy. Interestingly, the difference in accuracy disappeared when all tasks involving the investigation of relationship between two attributes became removed. This can be seen as an indicator that the quality of decision support heavily depends on the specific visualization implementation.

Amar and Stasko (2005) concluded that Kobsa's results indicate room for improvement of the standard operations but they also address that decision makers need more macro-level / statistical information about the data sets. They define the terms *worldview gap* and *rationale gap*. The *worldview gap* describes the discrepancy between the chosen representation and what is really needed to make a decision. The *rationale gap* refers to the black boxes in the reasoning process. Though data mining techniques often consist of rather less intuitive machine learning and reasoning approaches (e.g. support vector machines, neural networks) a decision maker needs to understand cause and effects. A decision based on a black box output is rather hard to explain and is perceived to be unethical.

Using a five step model they show how bridging the *worldview* and *rationale gaps* helps the user to accomplish his tasks with macro-level analytics.

Traditional visualization merely represents the raw data set and lacks the wider perspective. Amar and Stasko argue that this is not sufficient and that a representation should make its own limitations clear to its user. This is difficult and only seldom achieved by tools nowadays.

5.3.2 Applying Heuristics and Taxonomies

The goal of the prototype was to develop a new visualization to support users in some particular tasks. Three different visualizations have been implemented and can be checked against the taxonomy model of Roth and Mattis *(1990)*. Visualization for other tasks or data was beyond the scope of this work. The prototype visualizes other data only in text form.

The stacked bar charts of the trading volume are based on an algebraic dependency among complex data. The $total\,trading\,volume = \sum_{s \in S} trading\,volume_s$, meaning that the trading volume per time slot depends on the volumes of each share. Even these can be broken down again to the lot sizes of all transactions within the time slot. However, in the prototype a more general overview with a higher level of aggregation was preferred. Both aspects can be included in a single visualization by using a tree-map, which sets the main focus more on the lot size than on the total trading volume per day.

The price chart is a classical plot of scalars. These values have an inherent ordering by their time stamp. Furthermore, price can also be interpreted as coordinate by its inherent ordering. This means that the scatter plot is a valid display form, while e.g. a map view is not.

The most interesting point is the graph visualization. In the taxonomy of Roth and Mattis this visualization is called node-link-diagram. It is applicable for complex relationships among complex data types. They point out that a node-link-diagram shows very well *no-values*, but are ineffective when non-coverage is due to missing data. This means that if there is missing data, the absence of a link does not necessarily mean that there is no link in the real world. It is simply not in the data set. This is partly relevant for the filtered visualization presented in the previous chapter. Since the prototype is connected to the trading platform, there is no missing data. But to reduce the visual complexity, edges are omitted according to the user's filter settings. This may mislead the user. To overcome this problem, the user can open the filter for a selected node (see section 4.3).

These early visualization taxonomies are only partially helpful for evaluation purposes. These taxonomies are mapping data types and operations to visualizations. Newer taxonomies try to focus on the knowledge precepts of the tasks, which the user intends to solve with the visualization. The idea is to support the user in their knowledge deriving and decision making process. Applied to the use cases of the prototype, it is especially necessary in the fraud detection process. The support is twofold. First, the algorithms help to mark suspicious activities and accounts. Later on, the visualization brings the accounts into the market and social context. Still, one could think of more algorithms and more visualization. But the prototype is only a first approach towards integrated visual analysis tools for market operators.

According to the presented taxonomies, the prototype does use the adequat visualization techniques and supports the user to accomplish their tasks. What support users expect from such a tool is presented in the next section.

5.3.3 Fullfillment of Requirements

In the experiment, the users had to write down their key requirements for a market monitoring tool and stack rank them. After completing the tasks with the prototype they were asked about how well the prototype matched their expectations. Because of the small subset of observations (n=12) only a qualitative summary will be presented here. The most often mentioned requirement was the ease of use (seven times), followed by the clearly structured, holistic overview of the current situation (five times). Each of the following requirements have been stated at least three times: offering different levels of detail/drill down/filtering, highlighting outliers & detected fraudsters, automating the detection and optimizing the accuracy. Only one participant desired to have all cases presented which are suspicious such that he can filter them manually afterwards. Regarding the importance, most of these tasks have been ranked 2-3 on average with the only exception of automation ranked four. In the post experiment questions, the users could give grades between 1 (poorest) and 7 (best) and 0 for no answer/not applicable. It is a known phenomena that users tend to give better grades to systems that they have used on their own (compared to systems they have only seen pictures/videos of). Considering this, it is not surprising that the grades for the prototype are between 5 and 6.6 with an exception of 4.3 for the highlighting.

Most of the requirements mentioned by the participants are straight forward and are consistent with the requirements gathered from the literature in section 3.2.1. There is no way to guarantee that all possible requirements have been discovered. Many are taken from literature which only describe similar fields of applications such as money laundering and credit card fraud. But the way the user interacts with the application is very similar. The full list is given in table 5.5, now extended with the requirements fulfilled by the presented visualization.

How the prototype matches these requierments: The query interface is realized in various ways. The account inspector (c.f. 4.1.6), share inspector (c.f. 4.1.5), graph visualization (c.f. 3) and the account registration data dialog allow the user to narrow his search by time, share or user-name. Data from other sources are possible to incorporate: e.g. realized events such as results of a soccer match can be added and used to enrich the visualizations. Temporal links (i.e. transaction X was executed shortly after the account Y had been created) in the sense of Sentator et al. (1995) are not incorporated in the prototype. A first approach of temporal links in market places can be seen in the work presented by Maranzato (2010).

The algorithms are designed to allow continuous detection in the background while visualizing the market activity (see Chapter 2 on page 19). They treat every market participant equally in the sense that each algorithm does not take previous classifications of another

Table 5.5: Requirement list for a market visualization for fraud detection

	Senator et al.	Chau & Faloutsos	Chang et al.	this approach
Application				
query interface ‡	☑	☑	☑	☑
incorporate data from other sources†	☑			□
"temporal links"‡				
integrated and continuous detection†	☑			☑
periodic evaluations and update of the pattern†¶	☑	☑		
equal treatment of all market participants†	☑	□	☑	☑
Visualization				
weighted node positions				☑
drill down$			☑	☑
filtering$			☑	☑
(near) real-time interaction$			☑	☑
automatic layouts$			☑	☑
identification of "key" nodes$¶		☑	☑	☑
similar sub graphs$¶		☑	☑	
marking outliers¶		☑		☑
central spot of information¶		☑		☑
adequate visualization for different user groups†				
highly flexible, generic view concepts$			□	□
different visualization forms$			☑	☑
adaptable graphical parameters$			☑	☑

† Senator et al. $ Chang et al.

‡ Goldberg et al. ¶Chau and Faloutsos

algorithm or human into consideration. This decision was made on purpose since every algorithm has a certain false positive probability. If and account is treated by another algorithm as a fraudster when it has not yet been confirmed, the algorithm may produce bad results. Human classifications could be taken into consideration. However, the algorithms were designed not to depend on a preclassified (sub-)set. Though the algorithms detect a certain pattern within the graph, they do not perform a generic, similar subgraph detection in the sense of Chau et al (2006).

Key nodes can be identified easily since the modified graph layout displays them in the center of the visualization. Though the layout is completely automatic by the modified spring embedder algorithm, the user can exclude nodes and fix them to any desired location of the screen. This way, the visualization works as a central spot of information including three forms of visualizations (bar chart, line chart and graph visualization). From this central spot the user can drill down onto any desired data point. Graph visualization, line and bar charts are generic view concepts. The prototype limits the user to information which can be visualized by at least one of these views. The user can adapt the graph visualization by manipulating the layout algorithm parameters and the sliding window settings. Customizing the tool beyond this would have been out of scope for this work.

Summarizing the aspects above, it can be stated that most of the requirements have been met by the presented prototype consisting of the algorithms from chapter 2 on page 19 and visualization of chapter 3 on page 47.

5.3.4 Usability Evaluation by Participatory Observation

Participatory Observation was first applied in sociology. The first successful application in the field of usability and software engineering is mentioned by Potts (1993). It is an interdisciplinary study from Summerville et al. (1993) in which sociologists and software engineers work together in order to improve the usability of a flight database system for air traffic control. In sociology, this practice of observation in the field has a long tradition. In this setting, the observer sits next to the participant and tries to observe and evaluate the user interactions, problems, and time to complete the given tasks. Important influence factors are the habituation and adaptation during the participatory observation. Habituation refers to the long-term habits and familiarity with the environment. Adaption, on the other side, refers to the short term when new stimuli from the environment distract the participant. Depending on the expertise of the participant, this may have a different influence on the results and has to be taken into consideration by the observer.

In the HCI experiment described above the participants were all familiar with the location and three of them had previous experience in market surveillance or fraud detection. The observer measured the time the participants needed to classify each account or account pair for each task. The average time per task was for task one, 40:12 minutes (maximum:

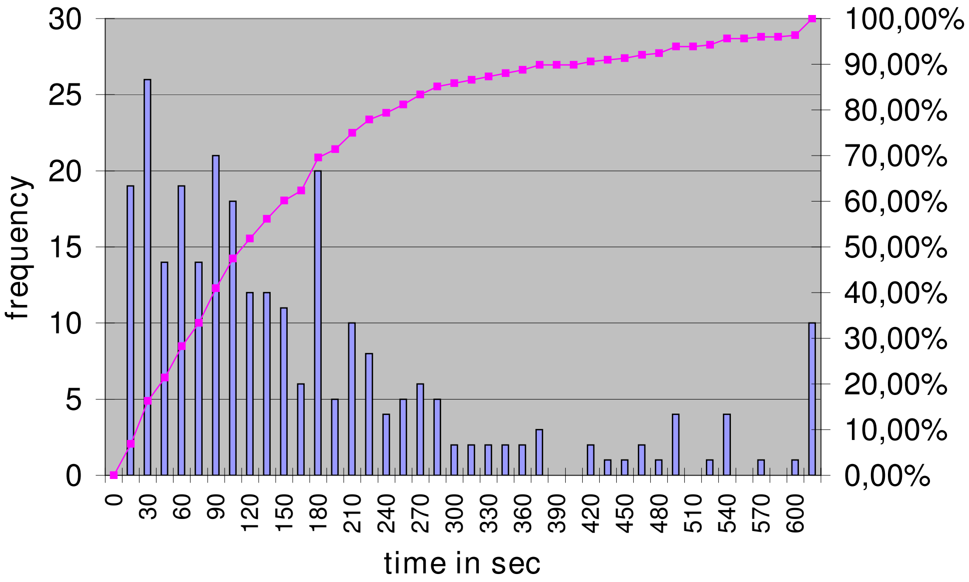

Figure 5.4: Histogram of the number of cases decided in less than x seconds

78:06 minutes; minimum 11:06 minutes) and for task two, 25:06 minutes (maximum 37:30 minuntes; minimum 10:36 minutes). Though faster, the participants percieved the task two as more difficult. The ranking scale was [1,2,3] with 1 for the most difficult task. They ranked task two with an average of 1.75, which was a little higher than task one with an average of 2.1.

After completing the difficulty ranking, the participants had to fill out a post experiment usability questionnaire[2] which was adapted from Lewis (1995). The questionnaire consisted of 20 questions regarding three areas: system usability, information quality, and interface quality. The participant can express how much he aggrees with the listed statement on a seven point Likert scale with an additional point for *N/A*. With an average of 5.6 (5.5 system usability, 5.5 information quality, and 5.7 interface quality) the participants were consistent with their qualitative statements in the requirement analysis (c.f. previous section).

Another interesting metric is the time until decision. It can be seen as the average time a user needs to gather enough insight to comfortably judge the case. The observer measured the time between the start of an investigation for a single account / account pair until the moment when the participant wrote his decision into the questionnaire. Figure 5.4 shows the histogram of the time. The x-axis is split in slices of 15 seconds, the left y-axis depicts the number of accounts judged within in this time for the bar charts, and the right y-axis the overall percentage of judged account for the line chart overlay. The participants were not aware of any time pressure. For example, there are

[2] The ranking and post experiment questionnaire were not part of the participatory observation.

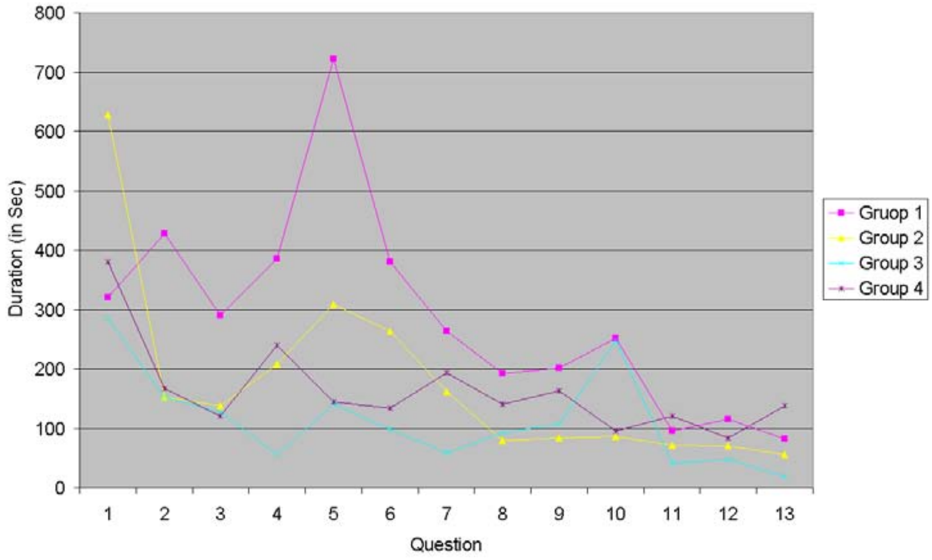

Figure 5.5: Group average times per account for task 1

10 instances in which the participants needed more than 10 minutes to decide. In one instance, a participant took even 22:30 minutes which causes the spike for question 5 in Figure 5.5. The line chart indicates 50% of the cases have been judged within two minutes (≤ 120 seconds) and 75% of all cases within 3:30 minutes.

These numbers have to be taken as upper bounds. The fact that only three persons had former experience in market surveillance and learning effects certainly influenced the time. The histogram is taken for all participants and all judgments. Because of learning effects, the evaluation of the first accounts took more time than for the following ones. This can be seen by comparing Figure 5.5 and Figure 5.6. Both figures have the same co-ordinate system: the x-axis describes the cases in the order the participants had to solve them, the y-axis the time in seconds. Though a line graph is not the correct presentation since each evaluation should be independent from the previous one, it reveals the trend over the 10 (13) cases. While in Figure 5.5 the average time for classification in each group descends; the graph of group three in Figure 5.6 is slightly rising. The average linear regression coefficient for task one is -20.2 while the average slope of the linear regression of task two is -6.5. This indicates that in the second task a certain amount of learning still took place, but less than in the first one. The coefficient of task two is also less robust with $\varnothing R^2$ of 0.15 compared to 0.40 for task one.

Overall, the participatory observation confirmed the principels of information visualization. Humans do perceive the task of identifying previously unknown patterns in a large amount of transaction data as more difficult than just verifiying the relationship between two accounts. While this is not surprising, the short time to decision indicates that the prototype supports the user well during both tasks. The learning curve indicates that most users were within less than an hour, including the 15 minutes tutorial video, al-

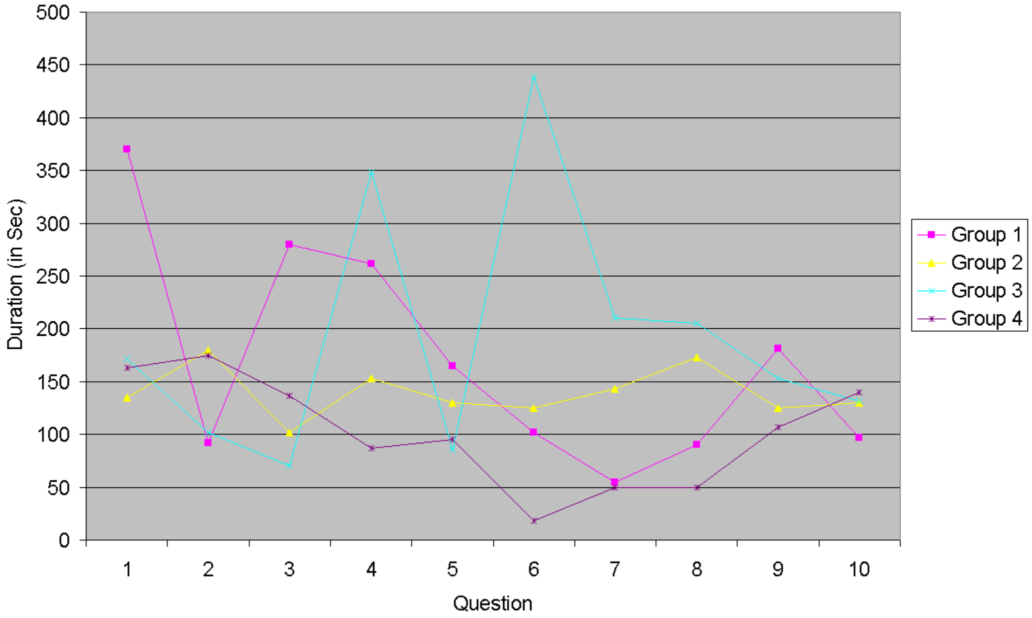

Figure 5.6: Group average times per account for task 2

ready very efficient in solving the given tasks. Even people without a market surveillance background achieved good results as the earlier presented interrater reliability shows.

5.3.5 Summary

The post-fieldstudy online survey has revealed that a significant amount of traders are tempted to try fraudulent strategies if the market were to be repeated again. Still, most of the people admitting fraud in the questionnaire had been found neither by hand nor by the algorithms. These were users who abandoned accounts with which they had backed the wrong horse and now wanted to have a second chance. From a prediction market perspective, this activity is double-edged. Since this behavior does not necesarrily harm the prediction quality if it happens purely sequential and always with best effort, it may even be supported by the trading platform. A functionality such as "close this account and create a new one" could be incorporated into a trading portal. The down side of such an approach is that it removes the pressure for participants to really process all available information and come up with a reasonable price. Instead, they could blindly bet on a share, simply to maximize their profit without even knowing anything about the market, the shares or the background information. Solving the problem remains for the market engineer of the futur prediction market platforms as it is beyond the scope of this work.

Furthermore, the accuracy of the algorithms has been objectified. With 88% accuracy, the algorithms slightly surpass the manual classification accuracy of 86%. Even better, the algorithms discover more accounts with higher trading volume than the manual classification. Besides the accuracy of the algorithms, the usability of the prototype itself has been tested and received an average of 5.6 on a 7 point Likert scale, which are very good results. Regarding the user expectations as well as the requirements gathered from the literature, the prototype fulfills almost all of them. This is confirmed by the steep learning curve of the participants. Within a short period of time, the participants were able to judge 75% of the cases in less than 3.5 minutes.

This shows that the graph based algorithms are a big step forward, since they warrent higher accuracy than manual spot checks on a real world data set and the visualization helps users to quicker separate fraudsters from innocent users.

6 Summary and Outlook

Today's electronic markets in general and play money prediction markets in particular are vulnerable to fraud and manipulation. There are many types of manipulations that can occur, ranging from the manipulation of reputation systems by using fake accounts, to price manipulation in order to increase profits and manipulation of the outcome on prediction markets. These three examples illustrate the variety of electronic market aspects which fraudsters target. In physical markets, processes are limited to the speed of action of the participants, however, electronic markets allow for a much faster pace. In order to keep track of transactions and ensure legal behavior, computer aided market surveillance systems are being employed. Scientists tend to use complex, supervised approaches, like neural networks, to detect fraud. These approaches depend on labeled data sets. In order to compare different suggested approaches, a common metric is necessary. However, the problems of missing agreement on a common metric, unpublished data sets, and closely following fraudsters, hinder further development of the field. Moreover, most industry solutions are unpublished.

Electronic markets —commercial, as well as academic prediction markets— need the customer to trust in the platform and fellow participants in order to be successful. Manipulation, and the resulting loss of reputation, is a serious threat to platform operators. If participants believe that others are using illegal methods, their willingness to commit fraud increases. This is shown in the post field study questionnaire of the participants of STOCCER, taken nine months after the end of the STOCCER main market. The study revealed that almost 10% of the questionnaire participants would commit fraud if STOCCER was repeated; this is twice the currently estimated amount of fraudsters.

Missing tools and insufficient staff are primary obstacles to monitoring the high transaction volumes of electronic markets. The high-level goal of the monitoring is quality assurance. One possible monitoring and fraud detection approach has been implemented as a prototype and is presented in the previous chapters.

6.1 Summary of Main Results

Traders are social entities trying to make a profit on a market. Market rules and platform restrictions usually prevent a single account from committing fraud or manipulation. However, many platforms are vulnerable to collusion. Furthermore, fraudsters tend to share ideas and techniques among friends or to simply copy what they observe on the market platform. Since collusion is a social process, the prototype presented in this work uses the transaction and user data to build the underlying social network. This network is mined for possible patterns of fraud.

Two graph based approaches to detect collusion in prediction markets have been presented. The design goals were to find simple and fast heuristics that are easy for market authorities to control and that achieve a high precision. As the previously mentioned disagreement on measurements indicates, precision is only one possible criterion. It was chosen following feedback from two employees of different market surveillance institutions who stated that a system which returns too many cases with poor precision will cause the staff to abandon the system or ignore the alerts.

The first approach (*Circular Trading Indicator*) tries to find repeated transactions between the same counter-parties, which is beneficial for only one of the two. The sensitivity analysis shows that the parameters allow for the exclusion of false positives with high certainty. The result set is shrunk to include only the absolutely obvious cases but increases the false negatives.

The second approach (*Prominent Edge Indicator*) is an example of a less accurate pattern which targets cases where one account is the major trading partner of another account. This trading concentration is measured by the Gini coefficient. Obviously, this can happen in absolutely legal and honorable transactions as well, for example, the first transaction a new account executes on the market. But besides this special case which can be pre-filtered, it selects accounts deviating from regular trading patterns by developing a preferred trading partner. Being less accurate, the false positive rate is 50%, much higher than the first approach. Many cases are just coincidence, for example, people might just happen to be interested in the same shares and trading at similar times. A sensitivity analysis of the different parameters showed that for this approach the precision could not be significantly improved.

Following the idea of finding the most relevant of the suspicious cases, a combination of the two approaches has been tested. The result is a stronger preselection without the necessity of calibrating any parameters as both indicators can be used in their least restrictive settings. Thereby achieving a higher than 90% precision on a real world data set. The importance of the low false-positve rate and tuning effort is crucial to find acceptance in the real business world as *Symantic* already underlined for Spam detection (Lochmaier 2009).

In terms of speed, the simpler second approach needs just 30 seconds for the complete replay of 32 graphs with a total of 17432 nodes and 47588 edges. The more complex indicator takes less than seven minutes.

The findings of Cortes et al. (2001) about fraudsters being closer to other fraudsters than normal users has been extended from telecommunications to electronic markets.

The test set includes the accounts tagged during the market run-time, plus the newly discovered accounts via the combination of the *Circular Trading* and the *Prominent Edge Indicator*. The shortest path from each account to the next fraudulent one has been measured. The hypothesis, that the length of the shortest path between fraudsters is longer than the shortest path between legitimate accounts and fraudsters, was tested using the Mann Whitney U Test. In the manually classified set, only 8 of the 32 markets had a p-value < 0.05 (see Section 2.2). After applying the indicators, the hypothesis H_0 can be rejected for 21 markets. The p-value for 6 of the remaining 11 markets is $<= 0.10$ and for the others dramatically decreased compared to the first test in the beginning of chapter 2. The interpretation of these results indicates that fraudsters were the main drivers in these markets since they had not been detected during the market operation. The only exception is the market of Switzerland where no fraudsters were detected by the algorithms and the values stayed the same.

In order to use the human abilities of pattern matching and analytical thinking, a visualization has been presented. The goal was to visualize the dynamics of a market over time in order to allow real time monitoring without overloading existing solutions with semantic complexity. Since most market surveillance teams are familiar with reading price and volume charts, these two displays were combined with a graph visualization of the transaction network between the different traders. This reveals correlations between different price or volume artifacts with different trading parties. Several enhancements to the standard spring embedder algorithm have been presented in order to enrich the graph visualization with more market related semantics. The run-time analysis shows that the modifications do not significantly increase the complexity (still in $O(n^2 \log n)$) and the empirical test indicates that the visualization accommodates common market activity very well within the dimensions of a common monitor.

The indicators and the monitoring component have been implemented in a prototype. This prototype has been successfully used on three other prediction markets so far and tested on a regular stock market data set. The usability of the prototype has been evaluated in an HCI experiment. In this experiment, 12 participants used the prototype to verify individual accounts (reflecting the scenario of a suspected fraud reported by a customer), verify account pairs (potential output of the indicators), and monitor the market during a limited time period. The results of the experiment show that the provided tools and visualizations were quite intuitive since the learning curve was quite steep. Moderate agreement was achieved regarding the classification quality of the untrained participants. They classified 50% of the cases in less than 2 minutes, and 75% of the cases within 3.5 minutes. This shows that the participants could gather the amount of data they needed to make a decision in a short period of time. In the HCI evaluation, the overall system quality was rated with 5.6 on a 7 point scale (5.5 system usability, 5.5 information quality, and 5.7 interface quality).

6.2 Limitations of the Approach

A major limitation during the development and evaluation of fraud detection methods is the unknown number of false negatives. False positives can be recognized rather easily. Even labeled data sets usually contain only the cases which have been revealed during previous investigations. Easy to classify or even pre-classified data sets could be generated via experiments. The scenario for the experiment would be two groups of traders, where one is technically and by incentives restricted to honest trading, while the other group is allowed and motivated to trick the system by all means. This method is still limited by the creativity of the participants and will most likely generate only small data sets. Generating data sets algorithmically calculates the false positive and false negative rates without uncertainty. However, it is questionable whether applying fraud detection mechanisms to a generated data set is a good approximation of real world scenarios. The consequence of using real world data sets is the unknown number of fraudsters, and therefore, the inability to determine precision and recall of a method. In the previous chapters, it has been assumed that the known set of fraudsters is the complete set of fraudsters.

Regarding the validation of the approach, it remains uncertain whether the classified fraudsters are really fraudsters. Some of the accounts have been excluded during market operations without complaint and some admitted fraud during the post field study questionnaire. But for all other accounts, there is only the judgement of the three human evaluators, where at least two of them were convinced that the accounts in question were fraudsters. Whether or not they really were fraudsters will probably remain uncertain forever.

The indicators are designed to analyze edges of the network. Therefore, they scale with the numberof edges m and nodes n of the graph. So far, the analyzed transaction graphs have been rather sparse. This is typical for most online markets. However, the graph might have a higher density on other platforms. For an online analysis the run-time complexity is $O(\frac{m}{n^2})$ for the circular trading indicator and $O(n)$ theoretical worst case for the prominent edge indicator. Because of the low density, the expected run-time is $o(\log n)$. These indicators are still suitable even for large graphs as the main factor influencing the complexity is the overall graph density. Since the so far observed real world market graphs are sparse with few dense clusters these indicators can easily be offered as a service to market operators. Since the fraud detection does not have to be a validation during transaction execution, it can be performed asynchroneus as a call to a service returning an alert. In case of higher transaction frequencies the indicators can easily be parallized in map-reduce fashion as long as no immediate neighboors are involved in the current set of transactions, such as $\overrightarrow{ab}$ and $\overrightarrow{bc}$. For the prominent edge indicator such transaction sequences need to be evaluated sequentially. Offering fraud detection as a service raises the concern whether it can be ensured that the clients are only real market operators instead of fraudsters who just want to test various strategies and whether they would get detected.

Since most of the edges are removed for the visualization, the goal of real-time monitoring depends on the computational power to calculate the force directed layout for all nodes. Though the run-time complexity is $O(n^2 \log n)$, it is too high for really large networks beyond 10,000 nodes. However, most markets can be split into a separate graph for each share. These graphs are usually much smaller than the overall registered number of accounts on the market platform. For example, on eBay, the group of bidders for one particular item. Though this graph might not yield enough evidence, it can easily be enriched by including transactions of previous co-occurances for each account pair.

Graph patterns in markets are still a relatively new field of research. New indicators for different patterns still have to be identified and implemented. Since fraudsters are closely following the new developments in the field, both auto-adjusting parameters and learning pattern detectors are desirable features for monitoring software

6.3 Outlook

Visualizations bear endless possibilities for modifications and extensions. An example of an extension is the encoding of further information into the shape of the nodes. Brandes et al. (2001) presented a graph visualization where an interval scale is encoded into the distortion of the node shape. This could be adapted for the market visualization in the following way: If a node is more of a liquidity provider than a requester, the shape might be more like a vertical ellipse, while if it is more of a requester, the shape is more like a horizontal ellipse. A balanced account would be depicted as a circle. Encoding further information into the graph visualization bears the subliminal danger of overloading the visualization. The simpler the visualization, the more intuitive is the resulting application following Eick's visualization mantra: overview first, details on demand.

Further enhancement of the indicator set can increase the versatility of the prototype. Examples of further indicators have already been presented in Section 2.7. Looking at the results of the infestations from the STOCCER market, the de-duplication of registration data yields a high potential for finding fraudsters on online platforms. Amateurs, who possess limited creativity in creating different, unlinked registration data, have been found more frequently than expected. Thus, linking accounts with similar registration data or the same pattern like "email=aaaaa@xyz.com" and "email=bbbbbb@xyz.com" would be a valuable extension of the existing set of indicators.

Looking at the requirements, there is one thing missing in all presented fraud detection tools: the *temporal links* from Goldberg and Senator (1995). They describe relations like "shortly after the account x was created, y made a huge sale to z". Linking all events with each other leads to exponential growth of the edge set with each occurring event.

Strong filtering will be necessary to remove edges between unrelated events. It remains an open question whether a graph is the best internal representation to mine these temporal relations.

Transaction data between users does not occur only on markets. Exchanged emails, bank transfers, and text messages from cell phones are a few other types of interactions which could be visualized using the same approach. Adding further types of transaction data could enhance the network formation. The visualization design is not limited to markets, but can be applied to other networks as well. It remains for future research to adapt the visualization for different types of interactions over time.

Social network analysis for prediction markets recently got more attention. *Intrade* released data sets to interested social network researchers in order to improve the accessibility, transparency, and validation of prediction markets (Intrade 2010). Social networking, in general, is getting more and more integrated into regular IT infrastructure (Strehlitz 2009; Gengler 2009). This opens the possibility to analyze even more transactions for audit purposes inside a company by using their social context.

The prototype presented here, has been used successfully to detect fraudsters in further markets and in a follow up research project (eix-market.de). According to Goldmann, the combination of our ego-driven, winner-takes-it-all mentality, as well as the pressure to deliver high growth rates, opportunities to commit fraud, and the culture of short-term earnings, all come together to foster the increase of fraud in our society (2010). Until we re-think our incentive systems and ethic values, the race between fraudsters and prosecutors will continue, making the use and further development of tools such as the proposed prototype necessary.

A Data Structures

Node Data Structure

```
class Node {
  // identifier, e.g. userid of the user-record in the DB
  long      id;
  // SUM[i](n_i * p_i) => the volume of all transactions
  double    totalVolume;
  // the timestep (millis since 1970) of the last transaction
  long      lastTransaction;
  ...
}
```

Edge Data Structure

```
class Edge {
  Node      source;          // => the seller
  Node      destination;     // => the buyer
  // SUM[i](n_i * p_i) => the volume of all transactions
  double    totalVolume;
  // the timestep (millis since 1970) of the last transaction
  long      lastTransaction;
}
```

By storing on each of the objects the last modification timestamp and aggregates of each of the attributes of interest, the control logic on top can easily highlight recent modifications and their impact in the network. The timestamp, for example, allows to reflect changes in the network in a replay mode. The aggregates can be mapped directly in the visualization to visual attributes like size, color, shape or speed. Otherwise they can be kept to be displayed for further investigation (*details on demand*).

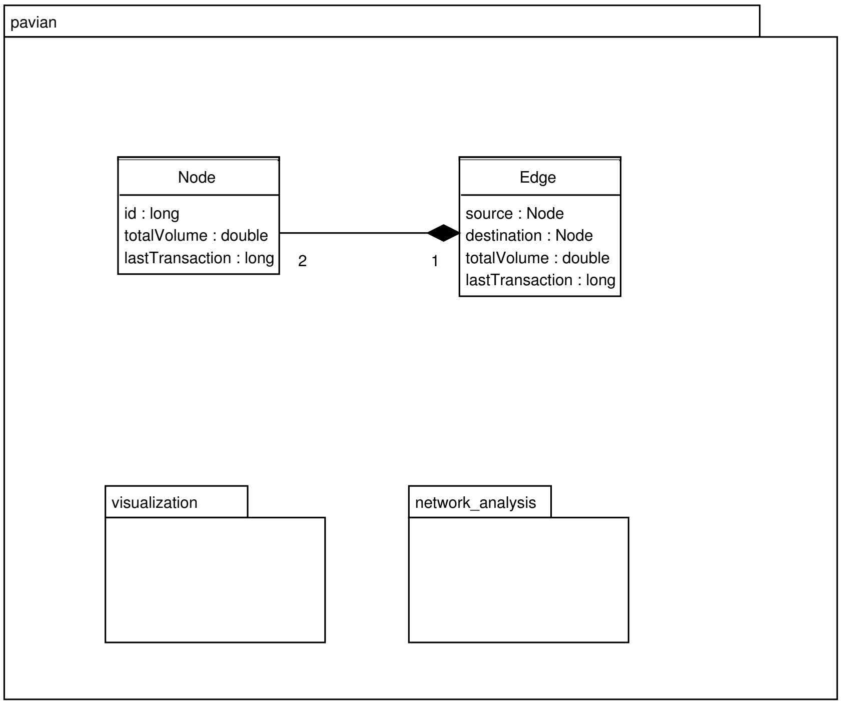

Figure A.1: Class Diagram

B Online Questionnaire

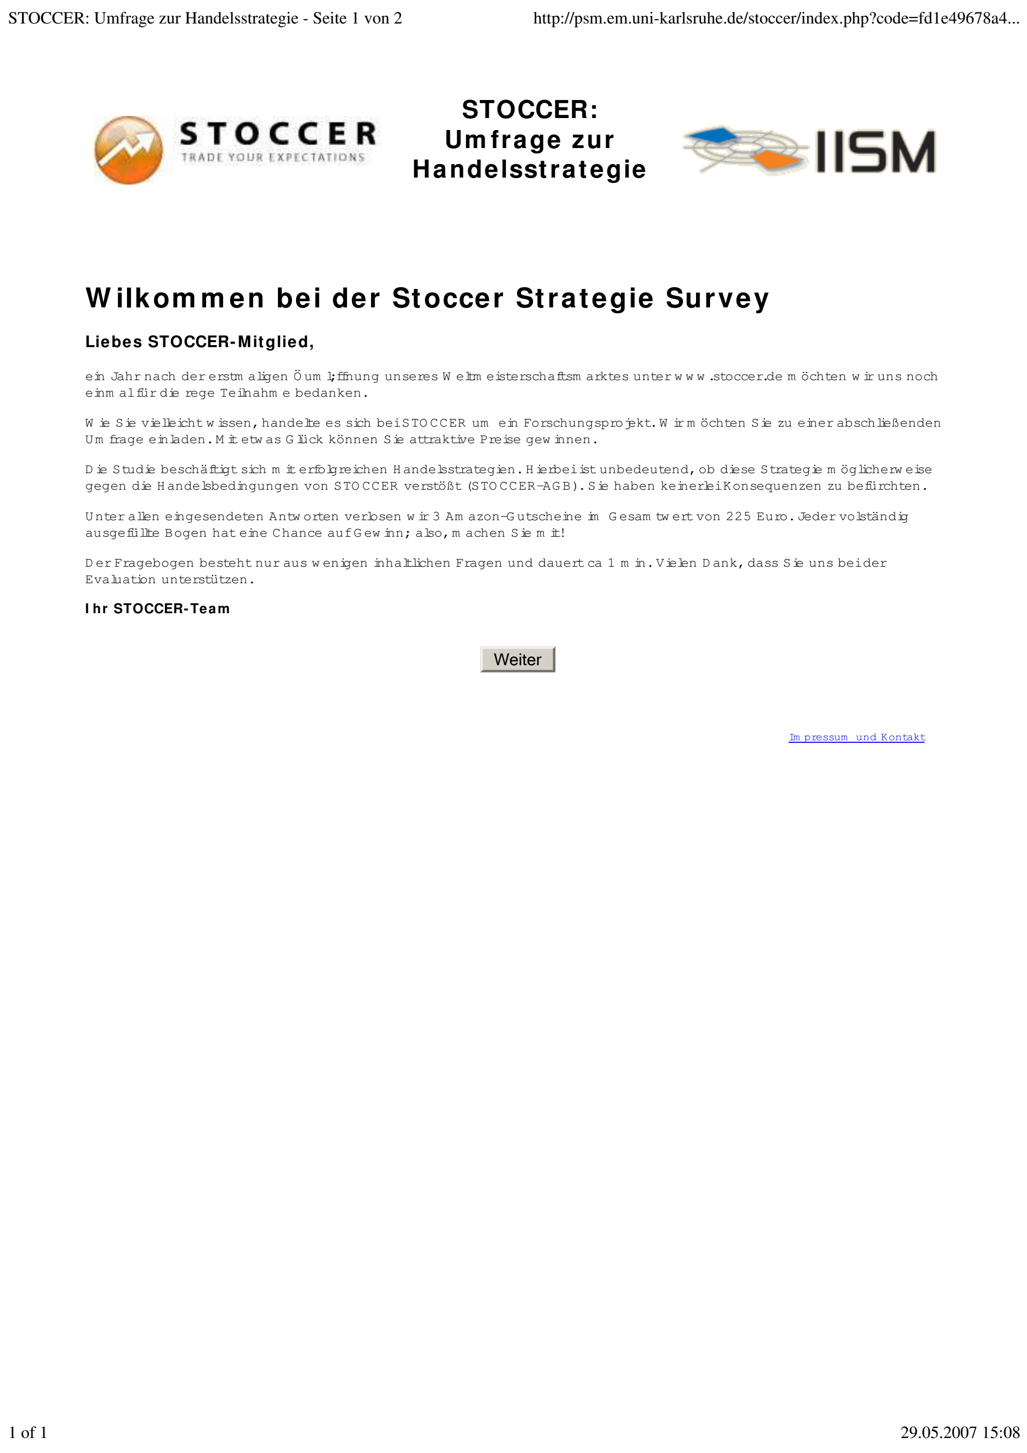

STOCCER: Umfrage zur Handelsstrategie - Seite 1 von 2 http://psm.em.uni-karlsruhe.de/stoccer/index.php?code=fd1e49678a4...

STOCCER: Umfrage zur Handelsstrategie

Wilkommen bei der Stoccer Strategie Survey

Liebes STOCCER-Mitglied,

ein Jahr nach der erstmaligen Öuml;ffnung unseres Weltmeisterschaftsmarktes unter www.stoccer.de möchten wir uns noch einmal für die rege Teilnahme bedanken.

Wie Sie vielleicht wissen, handelte es sich bei STOCCER um ein Forschungsprojekt. Wir möchten Sie zu einer abschließenden Umfrage einladen. Mit etwas Glück können Sie attraktive Preise gewinnen.

Die Studie beschäftigt sich mit erfolgreichen Handelsstrategien. Hierbei ist unbedeutend, ob diese Strategie möglicherweise gegen die Handelsbedingungen von STOCCER verstößt (STOCCER-AGB). Sie haben keinerlei Konsequenzen zu befürchten.

Unter allen eingesendeten Antworten verlosen wir 3 Amazon-Gutscheine im Gesamtwert von 225 Euro. Jeder vollständig ausgefüllte Bogen hat eine Chance auf Gewinn; also, machen Sie mit!

Der Fragebogen besteht nur aus wenigen inhaltlichen Fragen und dauert ca 1 min. Vielen Dank, dass Sie uns bei der Evaluation unterstützen.

Ihr STOCCER-Team

1 of 1 29.05.2007 15:08

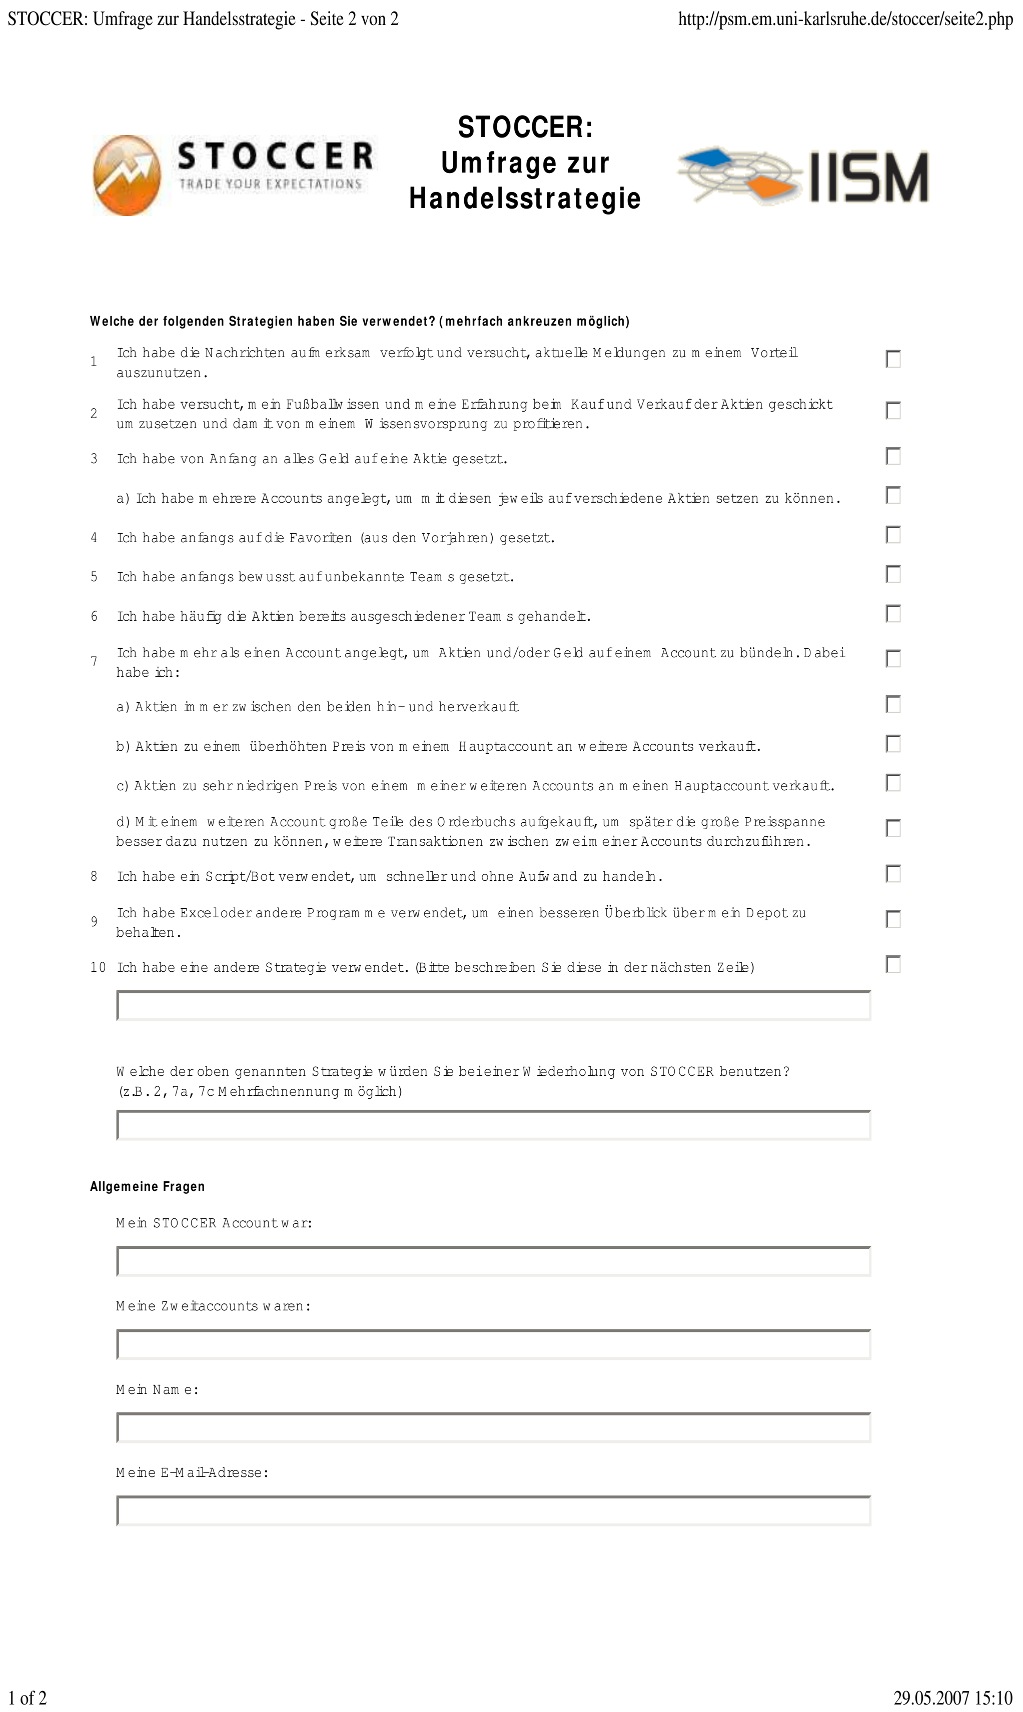

STOCCER: Umfrage zur Handelsstrategie - Seite 2 von 2 http://psm.em.uni-karlsruhe.de/stoccer/seite2.php

STOCCER: Umfrage zur Handelsstrategie

Welche der folgenden Strategien haben Sie verwendet? (mehrfach ankreuzen möglich)

1 Ich habe die Nachrichten aufmerksam verfolgt und versucht, aktuelle Meldungen zu meinem Vorteil auszunutzen. □

2 Ich habe versucht, mein Fußballwissen und meine Erfahrung beim Kauf und Verkauf der Aktien geschickt umzusetzen und damit von meinem Wissensvorsprung zu profitieren. □

3 Ich habe von Anfang an alles Geld auf eine Aktie gesetzt. □

a) Ich habe mehrere Accounts angelegt, um mit diesen jeweils auf verschiedene Aktien setzen zu können. □

4 Ich habe anfangs auf die Favoriten (aus den Vorjahren) gesetzt. □

5 Ich habe anfangs bewusst auf unbekannte Teams gesetzt. □

6 Ich habe häufig die Aktien bereits ausgeschiedener Teams gehandelt. □

7 Ich habe mehr als einen Account angelegt, um Aktien und/oder Geld auf einem Account zu bündeln. Dabei habe ich: □

a) Aktien immer zwischen den beiden hin- und herverkauft □

b) Aktien zu einem überhöhten Preis von meinem Hauptaccount an weitere Accounts verkauft. □

c) Aktien zu sehr niedrigen Preis von einem meiner weiteren Accounts an meinen Hauptaccount verkauft. □

d) Mit einem weiteren Account große Teile des Orderbuchs aufgekauft, um später die große Preisspanne besser dazu nutzen zu können, weitere Transaktionen zwischen zwei meiner Accounts durchzuführen. □

8 Ich habe ein Script/Bot verwendet, um schneller und ohne Aufwand zu handeln. □

9 Ich habe Excel oder andere Programme verwendet, um einen besseren Überblick über mein Depot zu behalten. □

10 Ich habe eine andere Strategie verwendet. (Bitte beschreiben Sie diese in der nächsten Zeile) □

Welche der oben genannten Strategie würden Sie bei einer Wiederholung von STOCCER benutzen? (z.B. 2, 7a, 7c Mehrfachnennung möglich)

Allgemeine Fragen

Mein STOCCER Account war:

Meine Zweitaccounts waren:

Mein Name:

Meine E-Mail-Adresse:

1 of 2 29.05.2007 15:10

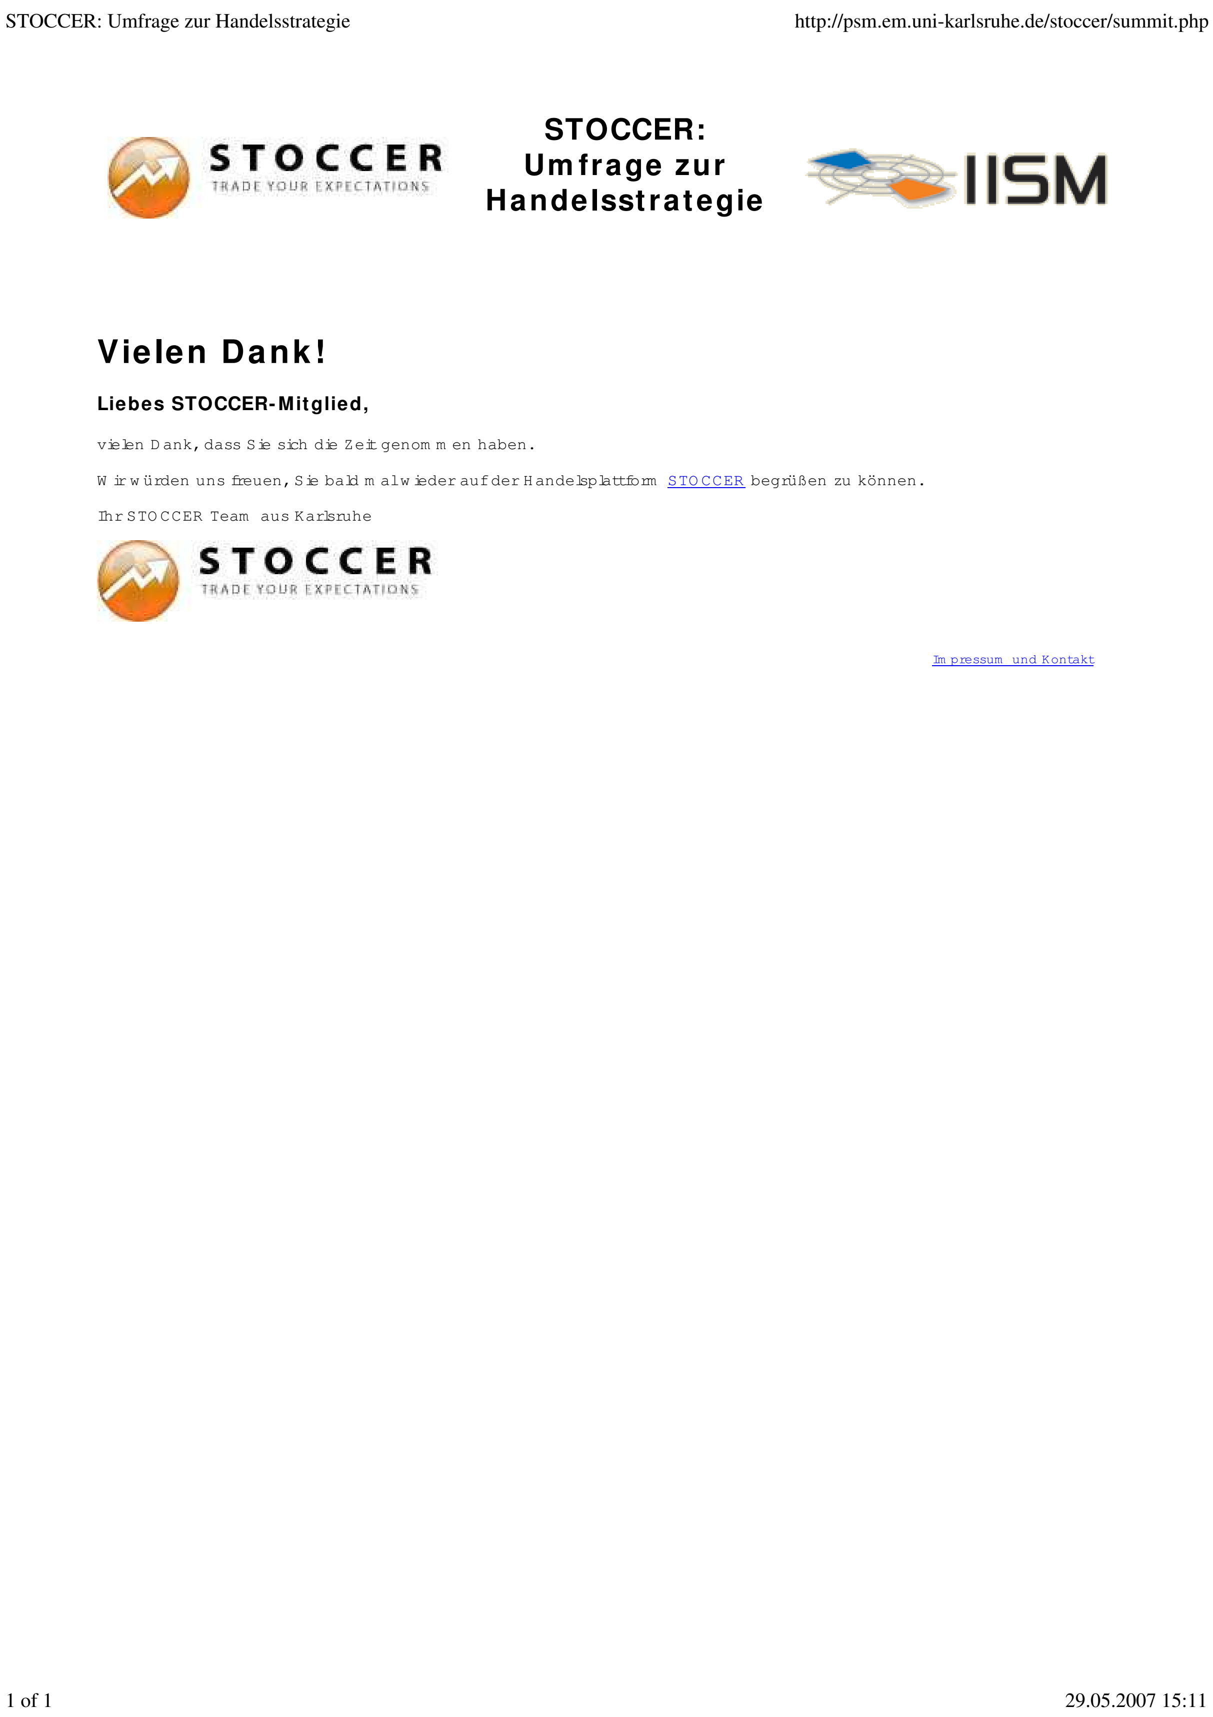

STOCCER: Umfrage zur Handelsstrategie http://psm.em.uni-karlsruhe.de/stoccer/summit.php

STOCCER: Umfrage zur Handelsstrategie

Vielen Dank!

Liebes STOCCER-Mitglied,

vielen Dank, dass Sie sich die Zeit genommen haben.

Wir würden uns freuen, Sie bald mal wieder auf der Handelsplattform STOCCER begrüßen zu können.

Ihr STOCCER Team aus Karlsruhe

Impressum und Kontakt

1 of 1 29.05.2007 15:11

C STOCCER Terms and Conditions

1. Participation in the forecasting market STOCCER is free of charge but requires a registration. During the registration process, the following data has to be provided correctly and completely: First name, last name, e-mail address, gender, year of birth, country of origin and postal address. By registering the user guarantees for the correctness of all declarations given in the registration process.
2. Eligibility for registration requires 18 years of age and the approval of the terms and conditions. No person is allowed to register more than once.
3. The user's chosen login name will be displayed on the website as part of the rankings and in other analyses.
4. The operator saves the data which is necessary for the operation and the reconstruction of the trading activity. The data will not be passed on to third parties. In this context, third parties are all organisations and persons who are not clearly named as operators of STOCCER.
5. The operator is not liable for server breakdowns and technical malfunctions.
6. The operator reserves the right to perform changes or supplementations to the existing offers without prior notice. The operator has the right to abort the exchange at any time.
7. The use of any kind of automated actions, e.g. scripts for manipulating the sport exchange, is strictly prohibited.
8. The operator may disqualify participants who violate these terms and conditions. In such cases prizes can be reclaimed ex post. Participants' accounts can temporarily be disabled if they are suspect of having manipulated the market.
9. Non-monetary prizes cannot be paid cash and are not transferable to other persons.
10. There is no legal claim of participation.
11. In case of the existence of winners with equal deposit value, fortune decides which one will be in the lottery.
12. Employees of the Institute of Information System and Management, University of Karlsruhe (TH), employees of the Chair for Business Administration, esp. Electronic Commerce, University of Frankfurt, as well as employees of organizations and enterprises related to STOCCER are eligible to participate, but not eligible to win prizes.

13. The publication of the forecasts created with this exchange for commercial purposes or in media requires the approval and, if applicable, the agreement of the operators.
14. The exchange is run under the law of the Federal Republic of Germany. Terms and conditions apply.

D Experiment Questionnaire

The questionnaire is printed with the same page breaks as the original one.

Introduction

Question 1

What do you expect of an information system for market monitoring and fraud detection? What requirements do you have? Please Rank them afterwards according their importance to you.

Expectation/Requirement	Rank

Please watch now the following videos in the given order:

1. 1_madevi_graph.wmv
2. 2_madevi_inspector.wmv
3. 3_madevi_fraud.wmv

List of Tasks:

1.) Checking account pairs

Scenario:

An algorithmic indicator provides a list of suspicious account pairs. The algorithm suggests that these pairs are partners/coalitions. One account is ruined in favour of the other typically trading large amount of shares back and forth between them in a short period of time. The algorithm stored the list on the hard disk in "task1.txt".

Task description:

You have to validate the list, classifying the accounts into normal users and fraudsters. If you are unsure, please mark the third option and write down a comment. Before you can start just open the application and select the market with the ID "9". Then go to the "analyze" menu and select the option "case manager". Select the option "Open fraud list..." from the toolbar and open the file "task1.txt". If you accidentally close the window, just open it again like described above. You can use further information windows from the view menu, if you need them. If you find more accounts just continue the list.

Id	Account	Id	Partner	Normal	Fraud	Unsure	Comment
102	Luckner	145	didi				
1845	tommy	249	benni39				
1488	Pia	1195	Maddin				
1398	Consti	975	Periastron				
567	Lady-W	480	e-union				
1308	dgut	780	Degauss				
519	vonderwand	791	LePompiste				
493	dududei	586	Kabeljau				
1019	falke	830	fpschebe				
1074	yb1898	1432	Uli				
1773	master	1746	king				
687	Juster	716	redondo				
1180	MB	780	Degauss				

2.) Checking accounts:

Scenario:

The call centre of your company filed several complains where participants were accusing other accounts to be violating the rules. The compiled list is handed over to you (since you are representing the market operation centre).

Task description:

Please open the Account Inspector from the "view" menu and investigate the given accounts. Again try to distinguish normal users from fraudsters. You may list some irregularities or difficulties you have in the comment field. If you are unsure about the true state of an account, mark the "unsure" column. If you find further accounts just continue the list.

Id	Partner	Normal	Fraud	Unsure	Comment
1874	hiphop				
1870	cocacola				
161	Spielmacher				
1847	Attila				
373	Liceu				
280	bullrider				
990	Studienrat				
1763	fruit				
728	diemerm				
292	fluxrope				

3.) Market Monitoring:

Replay the STOCCER World Championship Market (ID 9) from 2006-07-01 till 2006-07-05 in the market graph view and look whether you can observer any strange events. You have 15 minutes time.

Please write down your findings:

____

____

____

____

____

...

Final Questionnaire

Instructions

This questionnaire, which starts below, gives you an opportunity to tell us your reactions to the system you used. Your responses will help us understand what aspects of the system you are particularly concerned about and the aspects that satisfy you.

To as great a degree as possible, think about all the tasks that you have done with the system while you answer these questions.

Please read each statement and indicate how strongly you agree or disagree with the statement by circling a number on the scale. If a statement does not apply to you, circle N/A.

Thank you!

Question 1

How well did the application meet your initially defined requirements?

Expectation/Requirement	Very good				Very bad			N/A
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□
	□	□	□	□	□	□	□	□

Which was the most difficult task for you?

TASK	RANK	Why?
TASK 1: Checking Account Pairs		
TASK 2: Checking Accounts		
TASK 3: Monitoring		

Question 2 - System Usability

	Strongly agree				Strongly disagree			N/A
1. Overall, I am satisfied with how easy it is to use this system.	□	□	□	□	□	□	□	□
2. It was simple to use this system.	□	□	□	□	□	□	□	□
3. I could effectively complete the tasks and scenarios using this system.	□	□	□	□	□	□	□	□
4. I was able to complete the tasks and scenarios quickly using this system.	□	□	□	□	□	□	□	□
4. I was able to complete the tasks and scenarios quickly using this system.	□	□	□	□	□	□	□	□

	Strongly agree				Strongly disagree			N/A
6. I felt comfortable using this system.	□	□	□	□	□	□	□	□
7. It was easy to learn to use this system.	□	□	□	□	□	□	□	□
8. I believe I could become productive quickly using this system.	□	□	□	□	□	□	□	□
9. The system gave error messages that clearly told me how to fix problems.	□	□	□	□	□	□	□	□
10. Whenever I made a mistake using the system, I could recover easily and quickly.	□	□	□	□	□	□	□	□

Question 3 - Information Quality

	Strongly agree				Strongly disagree			N/A
11. The information (such as on-line help, on-screen messages and other documentation) provided with this system was clear.	□	□	□	□	□	□	□	□
12. It was easy to find the information I needed.	□	□	□	□	□	□	□	□
13. The information provided for the system was easy to understand.	□	□	□	□	□	□	□	□
14. The information was effective in helping me complete the tasks and scenarios.	□	□	□	□	□	□	□	□
15. The organization of information on the system screens was clear.	□	□	□	□	□	□	□	□

Question 4 - Interface

Note: The interface includes those items that you use to interact with the system. For example, some components of the interface are the keyboard, the mouse, the screens (including their use of graphics and language).

	Strongly agree				Strongly disagree			N/A
16. The interface of this system was pleasant.	□	□	□	□	□	□	□	□
17. I liked using the interface of this system.	□	□	□	□	□	□	□	□
18. This system has all the functions and capabilities I expect it to have.	□	□	□	□	□	□	□	□
19. Overall, I am satisfied with this system.	□	□	□	□	□	□	□	□

E Fraud Detection and Prevention Software

The following list is compiled from kdnuggets and related pages.

Toolname	Hersteller	Clients	Application	Technology
FraudNet 41st Parameter (2008)	41st Parameter	*Financial Institutions, Merchants*	fraud prevention and documentation	EZKeys, TimeDiff Linking, Data Spider, Sketch Match
ACI Retail Commerce Server ACI (2008a)	ACI	*Retail*	check acceptance	authorization system
ACI Proactive Risk Manager ACI (2008b)	ACI	*Wholesale*	card fraud detection, money laundering	user defined rules, neural networks
CCM (Continuous Control Monitoring) ACL (2008)	ACL	*Austrian Ministry of Finance*	Accounting Fraud	continuos audits
SAND Corporation (2008a)	Advanced Software Design Corporation	*Banks*	electronic check processing	encrypted authentication code
Fractals Alaric (2008)	Alaric	*Credit / Debit Card Companies*	card fraud	scoring model
Alcatel Fraud Management Systems Lucent (2008)	Alcatel Lucent	*Telecommunication*	Subscriber fraud	rule based system
ScorXPress ASA (2008)	ASA	*Credit / Debit Card Companies*	Credit card misuse	neural network

Toolname	**Hersteller**	**Clients**	**Application**	**Technology**
several Austinlogistics (2008)	austinlogistics	*Seven of the top 10 banks in the world**	Predictive analytic software solution	rule engines, prediction models
WatchDog Solutions (2008a)	Bassett Telecom Solutions	*Tele.ring Austria, Accenture Singapore*	telecommunications fraud	subscriber profiling, rating
iProtect, iPrevent, iComply Brighterion (2008)	brighterion	*Government, Banks, Insurance*	Intelligent Anti-Money Laundering Solution	Smart-Agents technology, Neural Networks, Case Based Reasoning
payment technologyCheckfree (2008)	carreker / checkfree	*Banks*	full range of payment processing solutions	scaned check image analysis
several ChoicePoint (2008)	ChoicePoint	*Government, Banks, Insurance*	customer profiling, insurance fraud, risk analysis	information integration
Magnify PATTERN:Detect Solutions (2008b)	ChoicePoint Claims Solutions	*Insurances*	insurance fraud	Analytics and Predictive Modeling, Accident Reports, Identity Matching, Smart Ordering, Compliance Watch List Screening, Carrier Identification, Visualization, Identity Verification, Case Management,Claims Monitoring

Toolname	**Hersteller**	**Clients**	**Application**	**Technology**
CyberSource Decision Manager CyberSource (2008)	cyberSource	*VISA*	transaction fraud	Fusion Scoring technology, Virtual Intelligence Risk Technology
LEADMiner technologies International (2008)	datamining-international	*US and partner governments*	structured and unstructured data analysis	expert systems, link analysis
FraudView Ectel (2008)	ectel	*Telecommunication*	subscriber fraud, identity fraud and many more	scoring model, rule engine, Link Visualization
equifax Equifax (2008)	equifax	*Credit / Risk Management*	information solutions	fraud experts do phone consulting
Falcon® Fraud Manager FairIsaac (2008)	FairIsaac	*Debit, credit, oil and retail card issuers*	detect and stop fraudulent transactions	rules engine
FinCAIS FINCEN (2008)	FINCEN	*US-Goverment*	money laundering, mortage fraud	link analysis
DebtIn4mer FML (2008)	FML	*service provider*	fraud in bad debt discovery solution	data mining
Fortent Fraud Management FORTENT (2008)	FORTENT	*Barclays, JPMorgan, The Bank of New York, RBS*	check fraud, money laundering	case based reasoning
Validis Route (2008)	Future Route	*accountants, financial directors and bookkeepers*	account validation and anomality detection	induction logic programming, genetic algorithms and information theory

Toolname	**Hersteller**	**Clients**	**Application**	**Technology**
HNC unveiled Autoadvisor(tm) INC. (2008)	HNC SOFTWARE INC.	*Insurance Companies*	insurance fraud	neutral networks
Analyst's Notebook I2inc (2008)	i2inc	*over 2,000 organizations worldwide*	visual investigative analysis software	link analysis
NORA IBM (2008)	IBM	*US Government*	homeland security	identity resolution
infoglide Infoglide (2008)	infoglide	*Airlines*	airline passenger screening	fuzzy matching over multiple databases
RiskShield inform GmbH (2008)	inform GmbH	*banking, insurance and telecommunication*	transaction fraud	real-time pattern recognition
InfoRate InfoScore (2008)	InfoScore	*Banks, Insurance*	customer profiling	scoring, data cleaning
Secure Science Corporation	IntelliFound	*E-Commerce*	credit card, identity fraud, etc.	monitoring
severalInterX (2008)	InterX	*several governement organisations*	distributed datasource pattern search	agent technology
F.A.L.S.T.A.F.F. Government (2008)	Italien Government	*Italien Ministery of Trade and Finance*	counterfeiting, customes	database and AI
Fraud Management Lavastorm (2008)	Lavastorm	*Telecommunication*	telephone fraud	knowledge-based system
VECTOR Fraud Solution Solutions (2008c)	Metavante Image Solutions	*Electronic Payment Companies*	check processing	rule selection and filtering

Toolname	**Hersteller**	**Clients**	**Application**	**Technology**
SONAR (ADS) NASD (2008)	NASD	*New York Stock Exchange*	NASDAQ	several AI techniques, such as data mining, natural language processing for text mining, intelligent software agents, rule-based inference, and knowledge-based data representation
Minotaur Technologies (2008a)	Natural Technologies (NT)	*Telecommunication and Financial Sector*	fraud management	neural predictive analytical models
PRISM eFraud Nestor (2008)	Nestor	*Internet Gaming, Lottery and Casino Customers*	detection and prevention	neural networks
NetMap (2008)	Netmap Analytics	*Law Enforcement and Intelligence Agencies*	crime investigation	map data finding common links
Neural Technologies Decider Technologies (2008b)	neural technologies	*Credit / Risk Management*	customer profiling	neural networks
Nt's Minotaur Fraud Management Solution Technologies (2008a)	neural technologies	*Telecommunication and Financial Sector*	subscription fraud, identity fraud	neural networks
NFC Sentinel NFC Global (2008)	NFC Global, LLC	*Banks*	customer profiling	data integration
Oscar Kilo's core product Kilo (2008b)	Oscar Kilo	*Banks, Insurance*	transaction fraud	pattern matching

Toolname	**Hersteller**	**Clients**	**Application**	**Technology**
DETECT Kilo (2008a)	Oscar Kilo	*All kinds of businesses*	credit card fraud, factoring fraud, insurance fraud, mobile phone fraud, etc.	risk engine
Red Shield ReD (Retail Decisions) (2008)	ReD (Retail Decisions)	*Ecommerce (b2c)*	card fraud prevention	neural networks, pattern recognition
RootStream Detect Stream (2008)	ROOT Stream	*Auditors*	accounting fraud	statistical techniques
FraudOffice SearchSpace (2008)	SearchSpace	*Ericsson*	telephone fraud	genetic algorithms, fuzzy logic, and neural network technology
MonITARS Searchspace (2008)	Searchspace	*London Stock Exchange, Bank of New York*	transaction fraud	genetic algorithms, fuzzy logic, and neural network technology
SMARTS Smartsgroup (2008)	smartsgroup	*Stock Exchanges, Regulator, Broker*	market surveillance	event processing system
AdvancedMiner StatConsulting (2008)	StatConsulting	*Telecommunication*	application fraud, billing fraud	scoring model
Nikira Subex (2008)	Subex	*Telecommunication*	fraud management system	rules-based alarms and pattern matching
risknet, redflag, smartauth Corporation (2008b)	The ai Corporation	*Credit / Debit Card Companies*	payment platforms	monitoring

Toolname	**Hersteller**	**Clients**	**Application**	**Technology**
piCARD Angency (2008)	The Modeling Angency	*Auditors*	credit card, procurement	pattern matching
LinkExplorer Tiburion (2008)	Tiburion, Inc.	*Public Safety and Justice Organisations*	public safety solutions	link analysis
Centrifuge Tildenwoods (2008)	tildenwoods	*US Government*	law enforcement	visual information analysis
VisuaLinks Inc. (2008b)	Visual Analytics Inc.	*US-Goverment*	homeland security	visual pattern discovery
Data Clarity® Suite Inc. (2008a)	Visual Analytics Inc.	*Commercial Business*	information sharing and pattern discovery solution	data mining, clustering, timeline analysis, social network analysis
WizRule WizSoft (2008a)	WizSoft	*Insurance, Auditors*	database analysis	rule engine, likelyhood estimation
WizWhy WizSoft (2008b)	WizSoft	*Insurance, Auditors*	database analysis	rule generator
Xtract Autoscore Xtract (2008)	Xtract	*Telecommunication*	customer relation analysis	neural networks
Fraud Prevention SmartSystem Beck (2008)	Beck	*Telecommunication*	billing fraud	finger printing, learning, rule engine
Data Mining in Identity Crime Prevention Universities (2008)	Several Universities	*Banks, Credit Card Companies*	white-collar crime	data mining, graph mining

Bibliography

The currency and foreign transactions reporting act, September 2000. URL http://www.occ.treas.gov/handbook/bsa.pdf. 31 USC Sections 5311-5330 and 12 USC Sections 1818(s), 1829(b), and 1951-1959.

41st Parameter. Fraudnet, 2008. http://www.the41.com/site/index.html.

D. Abbott, P. Matkovsky, and J. Elder. An Evaluation of High-End Data Mining Tools for Fraud Detection. In *Proc. of IEEE SMC98*, 1998.

ACFE. Report to the nation on occupational fraud & abuse. Technical report, Association of Certified Fraud Examiners, 2006.

ACI. Aci retail commerce server, 2008a. http://www.aciworldwide.com/frauddetectionsoftware/.

ACI. Aci proactive risk manager, 2008b. http://www.aciworldwide.com/frauddetectionsoftware/.

ACL. Ccm (continuous control monitoring), 2008. http://www.acl.com/solutions/fraud_detection.aspx.

L. v. Ahn, M. Blum, N. J. Hopper, and J. Langford. Captcha: Using hard AI problems for security. In *Eurocrypt 2004*, 2004a.

L. v. Ahn, M. Blum, and J. Langford. Telling humans and computers appart automatically. *Communications of the ACM*, 47(2):57–61, 2004b.

Alaric. Fractals, 2008. http://www.alaric.com/public/products/fractals.

F. Allen and D. Gale. Stock-price manipulation. *Rev Financ Stud*, 5(3):503–529, 1992.

J. Amar, R.A.; Stasko. Knowledge precepts for design and evaluation of information visualizations. *Transactions on Visualization and Computer Graphics*, 11(4):432–442, July-Aug. 2005. ISSN 1077-2626. doi: 10.1109/TVCG.2005.63.

N. Analytics. Netmap, 2008. http://www.netmap.com.au/crime.html.

K. V. Andersen, S. Elliot, P. M. C. Swatman, E. Trauth, and N. Bjorn-Andersen, editors. *Seeking success in ebusiness*. IFIP Advances in Information and Communication Technology. Springer, 2003.

T. M. Angency. picard intelligent procurement card monitoring system, 2008. http://www.the-modeling-agency.com/solutions/picard.html.

D. Archambault, T. Munzner, and D. Auber. Topolayout: Multilevel graph layout by topological features. *IEEE Transactions on Visualization and Computer Graphics*, 13 (2):305–317, 2007.

ASA. Scorxpress, 2008. http://www.asacorp.com/?&tc=true.

E. August Bequai. Organized crime goes cyber. *Computers & Security*, 20:475–478, 2001.

Austinlogistics. Analytic modeling, 2008. http://www.austinlogistics.com/products/analyticmodeling.php.

P. Bachmann. *Die Analytische Zahlentheorie*. B.G. Teubner, Leipzig, 2 edition, 1894.

J. Barnes and P. Hut. A hierarchical o(n log n) force-calculation algorithm. *Nature*, 324: 446–449, 1986. doi: 10.1038/324446a0. URL http://www.nature.com/nature/journal/v324/n6096/abs/324446a0.html.

Beck. Fraud prevention smartsystem, 2008. http://www.beckcomputers.com/aboutFPSSFrame.html.

E. Belhadji, G. Dionne, and F. Tarkhani. A Model for the Detection of Insurance Fraud. *The Geneva Papers on Risk and Insurance*, 25(4):pp. 517–538, 2000.

Bissantz & Company GmbH. Deltamaster. online, January 2008. URL http://www.bissantz.de/loesungen/.

M. Blume and C. Weinhardt. Information visualization in markets - the stock diamond. In *Human Computer Interaction*, 2008.

R. Bobrik, M. Reichert, and T. Bauer. Requirements for the visualization of system-spanning business processes. In *DEXA Workshops*, pages 948–954. IEEE Computer Society, 2005. ISBN 0-7695-2424-9.

P. J. G. Bohm and J. Sonnegård. Political stock markets and unreliable polls. *Scandinavian Journal of Economics*, 101(2):205–222, June 1999.

R. J. Bolton and D. J. Hand. Statistical fraud detection: A review. *Statistical Science*, 17(3):235–255, 2002. (with discussion).

U. Brandes. *Drawing Graphs*, chapter Drawing on Physical Analogies, pages 71–86. 2001.

U. Brandes, J. Raab, and D. Wagner. Exploratory network visualization: Simultaneous display of actor status and connection. *Journal of Social Structure*, 2(4):1–28, 2001.

Brighterion. iprotect, iprevent, icomply, 2008. http://www.brighterion.com/.

P. Brockett, R. Derrig, L. Golden, A. Levine, and M. Alpert. Fraud Classification using Principal Component Analysis of RIDITs. *The Journal of Risk and Insurance*, 69(3): pp. 341–371, 2002.

C. F. Camerer. Can asset markets be manipulated? a field experiment with racetrack betting. *Journal of Political Economy*, 106(3):457–481, 1998.

S. J. S. J. Cannan and G. A. M. Otten. *SQL — the standard handbook: based on the new SQL standard (ISO 9075:1922(E))*. MCGRAW-HILL, 1993. ISBN 0-07-707664-8.

M. Chalmers. Pearls, swines and sow's ears: Interface research inside a multinational bank. In *BNCOD 14: Proceedings of the 14th British National Conference on Databases*, pages 222–229. Springer-Verlag, 1996. ISBN 3-540-61442-7.

R. Chang, M. Ghoniem, R. Kosara, W. Ribarsky, J. Yang, E. Suma, C. Ziemkiewicz, D. Kern, and A. Sudjianto. Wirevis: Visualization of categorical, time-varying data from financial transactions. In *Proceedings of the IEEE VAST Symposium 2007*, pages 155–162, Sacramento, USA, October 30 - November 1 2007. IEEE.

D. H. Chau, S. Pandit, and C. Faloutsos. Detecting fraudulent personalities in networks of online auctioneer. In *Proceedings of PKDD 2006*, pages 103–114, Berlink Germany, 2006.

C. . Checkfree. payment technology, 2008. http://www.carreker.com/main/solutions/payments/payments_ove.htm.

A. Chen, P. Tsai, and J. L. Koh. Identifying Object Isomerism in Multidatabase Systems. *Distributed and Parallel Databases*, 4(2):143–168, 1996.

ChoicePoint. several, 2008. http://www.choicepoint.com/business/index.html.

C. Cook, F. Heath, and R. L. Thompson. A meta-analysis of response rates in web- or internet-based survey. *Educational and Psychological Measurement*, 60:821–836, 6 2000. URL http://epm.sagepub.com/cgi/content/abstract/60/6/821.

I. Coral8. Coral8. online, January 2008. URL http://www.coral8.com.

A. S. D. Corporation. Sand, 2008a. http://www.asdc.com/Sand/sand.html.

T. A. Corporation. risknet, redflag, smartauth, 2008b. http://www.aicorporation.com/.

C. Cortes, D. Pregibon, and C. Volinsky. Communities of interest. In *Proc. of IDA2001*, pages pp105–114, 2001.

C. Cortes, D. Pregibon, and C. Volinsky. Computational methods for dynamic graphs. *Journal of Computational and Graphical Statistics*, 12(4):pp. 950–970, 2003.

W. L. Cukier, E. J. Nesselroth, and S. Cody. Genre, narrative and the "nigerian letter" in electronic mail. In *System Sciences, 2007. HICSS 2007. 40th Annual Hawaii International Conference on*, pages 70–70, Jan. 2007. doi: 10.1109/HICSS.2007.238.

C. Culpepper. The day the red flag went up. *LA Times*, page D. 1, 12th of August 2007. Betfair's 'fraud team' could quickly see that the betting pattern was strange for the Davydenko-Arguello match at an obscure men's tennis tournament in Poland.

CyberSource. Cybersource decision manager., 2008. http://www.cybersource.com/products_and_services/risk_management/fraud_screening/.

C. Czernohous, W. Fichtner, D. Veit, and C. Weinhardt. Management decision support using long-term market simulation. *Inf. Syst. E-Business Management*, 1(4):405–423, 2003.

N. Dalvi, P. Domingos, Mausam, S. Sanghai, and D. Verma. Adversarial classification. In *Proceedings of SIGKDD04*, pages 99–108, 2004. URL http://portal.acm.org/toc.cfm?id=1014052.

C. Davidson. What Your Database Hides Away. *New Scientist*, 1855:28–31, 9 January 1993.

R. C. Dehmel. Fast hierarchical methods for the n-body problem, part 1. online, April 1996. URL http://www.cs.berkeley.edu/~demmel/cs267/lecture26/lecture26.html. (visited on 2007-11-26).

E. Deutskens, K. de Ruyter, M. Wetzels, and P. Oosterveld. Response rate and response quality of internet-based surveys: An experimental study. *Marketing Letters*, 15(1): 21–36, Februar 2004. doi: 10.1023/B:MARK.0000021968.86465.00. URL http://www.springerlink.com/content/p44j16m17v630624/.

D. A. Dillman. *Mail and Internet Surveys: The Tailored Design Method*. John Wiley & Sons, 2nd edition, 2000.

S. Dimitrov and R. Sami. Non-myopic strategies in prediction markets. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 200–209, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-169-9. doi: http://doi.acm.org/10.1145/1386790.1386823.

C. Domingos, R. Gavalda, and O. Watanabe. Adaptive Sampling Methods for Scaling Up Knowledge Discovery Algorithms. *Data Mining and Knowledge Discovery*, 6(2): pp. 131–152, 2002.

J. Drape. Web site puts focus on the fix in sports bets. online, May, 25 2008.

P. Eades. A heuristic for graph drawing. In *Congressus Numerantium*, pages 149–160, 1984.

Ectel. Fraudview, 2008. http://www.ectel.com/content.aspx?id=97.

M. Elfeky, V. Verykios, and A. Elmagarmid. Tailor: a record linkage toolbox. *Proceedings. 18th International Conference on Data Engineering*, pages 17–28, 2002. doi: 10.1109/ICDE.2002.994694.

Equifax. equifax, 2008. http://www.equifax.com/home/.

FairIsaac. Falcon® fraud manager, 2008. http://www.fairisaac.com/fic/en/product-service/product-index/falcon-fraud-manager/.

E. F. Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25(2):383–417, 1970.

FINCEN. Fincais, 2008. `http://www.fincen.gov/MortgageLoanFraud.pdf`.

U. Fischbacher. z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, June 2007.

J. L. Fleiss. Measuring nominal scale agreement among many raters. *Psychological Bulletin*, 76(5):378–382, 1971.

FML. Debtin4mer, 2008. `http://www.fraudmanagement.com/in4mersheets/DIdatasheet2p11.pdf`.

R. Forsythe, F. Nelson, G. R. Neumann, and J. Wright. Anatomy of an experimental political stock market. *The American Economic Review*, 82(5):1142–1161, December 1992.

FORTENT. Fortent fraud management, 2008. `http://www.fortent.com/solutions/risk-compliance-technology.php`.

T. M. Fruchterman and E. M. Reingold. Graph drawing by force-directed placement. *Software-Practice and Experience*, 21(11):1129–1164, 1991.

M. R. Garey and D. S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman, New York, 1979.

B. Gengler. Das social web wird unverzichtbar. *Computer Zeitung*, page 13, 13th, July 2009.

N. Gershon and S. G. Eick. Information visualization applications in the real world - business visualization applications. *Information Visualization Business Notes*, pages 66–71, 1997.

A. Geyer-Schulz, S. Luckner, J. Schröder, B. Skiera, C. Slamka, and C. Weinhardt. Empirical evaluation of call auctions in prediction markets. Working paper, Universität Karlsruhe (TH), 2007.

C. Gini. On the measure of concentration with special reference to income and wealth. In *Cowels Commission Research Conference on Economics and Statistics*, pages 73–80, Springs, Colorado, 1936. Colorado College.

H. G. Goldberg and T. E. Senator. Restructuring Databases for Knowledge Discovery by Consolidation and Link Formation. In *Proceedings of the First International Conference on Knowledge Discovery and Data Mining*, Montreal, August 1995.

R. W. H. Goldberg, Henry Wong. Restructuring transactional data for link analysis in the fincen ai system. In D. Jensen and H. Goldberg, editors, *Artificial Intelligence and Link analysis*, 1998.

P. Goldmann. *Fraud in the Markets: Why It Happens and How to Fight It*. John Wiley & Sons, April 2010.

I. Government. F.a.l.s.t.a.f.f., 2008. http://www.agenziadogane.it/wps/wcm/connect/ee/HomePageEn/Falstaff/.

J. Hansen, C. Schmidt, and M. Strobel. Manipulation in political stock markets - preconditions and evidence. *Applied Economics Letters Applied Economics Letters*, 11(7): 459–463, June 2004. URL http://edoc.hu-berlin.de/series/sfb-373-papers/2001-61/PDF/61.pdf</.

R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.

R. Hanson and R. Oprea. Manipulators increase information market accuracy. Working-paper, 2004. URL http://econ.ucsc.edu/faculty/roprea/manipTH.pdf.

R. Hanson, R. Oprea, and D. Porter. Information aggregation and manipulation in an experimental market. *Journal of Economic Behavior and Organization*, 60(1):449–459, August 2006. URL http://econ.ucsc.edu/faculty/roprea/manipEX1.pdf.

M. C. Hao, D. A. Keim, U. Dayal, and J. Schneidewind. Business process impact visualization and anomaly detection. *Information Visualization*, 5(1):15–27, 2006. ISSN 1473-8716. doi: http://dx.doi.org/10.1057/palgrave.ivs.9500115.

F. Hayek. The use of knowledge in society. *American Economic Review*, 35(4):519–530, 1945.

J. Heer and D. Boyd. Vizster: Visualizing online social networks. *Proceedings of InfoVis 2005*, 2005.

M. A. Hernandez and S. J. Stolfo. The merge/purge problem for large databases. In *SIGMOD '95: Proceedings of the 1995 ACM SIGMOD international conference on Management of data*, pages 127–138, New York, NY, USA, 1995. ACM. ISBN 0-89791-731-6. doi: http://doi.acm.org/10.1145/223784.223807.

C. Hibbert. Zocalo: An open-source platform for deploying prediction markets. online, 2005. URL http://zocalo.sourceforge.net/. Commerce.net.

T. P. Hill. A statistical derivation of the significant digit law. *Statist.Science*, 10:354–363, 1995.

V. Hodge and J. Austin. A Survey of Outlier Detection Methodologies. *Artificial Intelligence Review*, 22(2):pp. 85–126, 2004.

I2inc. Analyst's notebook, 2008. http://www.i2inc.com/Products/Analysts_Notebook/default.asp.

IBM. Nora (non-obvious relationship awareness), 2008. http://radar.oreilly.com/archives/2005/04/non_obvious_rel.html.

H. S. INC. Hnc unveiled autoadvisor(tm), 2008. http://www.leavcom.com/hm_hnc.htm.

V. A. Inc. Data clarity® suite, 2008a. http://www.visualanalytics.com/Products/dcs/index.cfm.

V. A. Inc. Visualinks, 2008b. http://www.visualanalytics.com/.

Infoglide. infoglide, 2008. http://www.infoglide.com/.

inform GmbH. Riskshield, 2008. http://www.riskshield.de/.

InfoScore. Inforate, 2008. http://www.infoscore.de/de/dienstleistungen/informationsmanagement/lsungen/plattformen/index_20060329173808.html.

Innovations GmbH. mlds (money laundering detection system). online, January 2008. URL http://www.mlds.info/02_produkte/mlds.html.

A. Inokuchi, T. Washio, and H. Motoda. Complete mining of frequent patterns from graphs: Mining graph data. *Machine Learning*, 50:321–354, 2003. URL http://www.springerlink.com/content/txh58706lm053121/fulltext.pdf.

S. Institute. Using data mining techniques for fraud detection: A best practices approach to government technology solutions, whitepapers. Technical report, http://www.sas.com, 1996.

D. International. Leadminer technologies, 2008. http://www.dataminginternational.com/.

InterX. several, 2008. http://www.inferx.com/products.htm.

Intrade. Press, research & testimonials, 2010. URL https://www.intrade.com/aav2/press/index.html.

ISO 9241. Guidance on usability standards. Technical report, ISO, 1996. URL http://www.iso.ch/iso/en/CatalogueListPage.CatalogueList?ICS1=13&ICS2=180.

M. D. Kaplowitz, T. D. Hadlock, and R. Levine. A comparison of web and email survey response rates. *Public Opinion Quaterly*, 68(1):94–101, 2004. doi: 10.1093/poq/nfh006.

O. Kilo. Detect, 2008a. http://www.oscarkilo.net/products/.

O. Kilo. Oscar kilo's core product, 2008b. http://www.oscarkilo.net/products/.

H.-C. Kim, S. Pang, H.-M. Je, D. Kim, and S.-Y. Bang. Constructing support vector machine ensemble. *Pattern Recognition*, 36:2757–2767, 2003.

K. Kiviluoto and P. Bergius. Analyzing financial statements with the self-organizing map. In *in Proc. WSOM 97, Workshop Self-Organizing Maps*, pages 362–367, 1997.

J. Kleinberg. Authoritative sources in a hyperlinked environment. In J. Kleinberg, editor, *Proceedings 9th ACM/SIAM Symposium on Discrete Algorithms*, 1998.

D. Knuth. *The Art of Computer Programming, Volume 1: Fundamental Algorithms*. Addison-Wesley, 3 edition, 1997.

A. Kobsa. An empirical comparison of three commercial information visualization systems. *Information Visualization, 2001. INFOVIS 2001. IEEE Symposium on*, pages 123–130, 2001. ISSN 1522-404X.

K. Kolitz and C. Weinhardt. MES - ein experimentalsystem zur untersuchung elektronischer märkte. In *Service Oriented Electronic Commerce*, pages 103–118, 2006.

Y. Kou, C.-T. Lu, S. Sirwongwattana, and Y.-P. Huang. Survey of fraud detection techniques. In *Proceedings of the 2004 International Conference on Networking, Sensing, and Control*, pages 749–754, 2004.

A. S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985.

A. S. Kyle. Informed speculation with imperfect competition. *The Review of Economic Studies*, 56(3):317–355, 1989.

E. Landau. *Handbuch der Lehre von der Verteilung der Primzahlen*, volume 2. B.G.Teubner, Leipzig, 1909.

S. M. Lang and P. C. Lockemann. *Datenbankeneinsatz*. Springer, 1995.

Lavastorm. Fraud management, 2008. http://www.lavastorm.com/.

W. Lee, S. J. Stolfo, and K. W. Mok. A data mining framework for building intrusion detection models. *sp*, 00:0120, 1999. ISSN 1540-7993. doi: http://doi.ieeecomputersociety.org/10.1109/SECPRI.1999.766909.

J. R. Lewis. IBM computer usability satisfaction questionnaires: psychometric evaluation and instructions for use. *International Journal of Human-Computer Interaction*, 7(1):57–78, Jan-March 1995. URL http://portal.acm.org/citation.cfm?id=204770.204774.

Y. Livnat, J. Agutter, S. Moon, R. F. Erbacher, and S. Foresti. A visualization paradigm for network intrusion detection. In *Information Assurance Workshop, 2005. IAW '05*, pages 92–99. IEEE, June 2005a. doi: 10.1109/IAW.2005.1495939. URL http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=1495939. Proceedings from the Sixth Annual IEEE SMC.

Y. Livnat, J. Agutter, S. Moon, and S. Foresti. Visual correlation for situational awareness. In IEEE, editor, *IEEE Symposium on Information Visualization*, pages 95–102. IEEE, 2005b.

L. Lochmaier. Algorithmen gegen spam kämpfen mit fehlalarmen. *Computer Zeitung*, (27):12, 29th, June 2009.

LottStV. Staatsvertrag zum Lotteriewesen in Deutschland. Bundesgesetzblatt, Dezember 2003.

A. Lucent. Alcatel fraud management systems, 2008. `http://www1.alcatel-lucent.com/doctypes/opgproductbrochure/pdf/FMSV1.pdf`.

S. Luckner. *Predictive Power or Markets*. PhD thesis, Fakultät für Wirtschaftswissenschaften, Universität Karlsruhe (TH), 2008.

S. Luckner, J. Schröder, and C. Slamka. On the forecast accuracy of sports prediction markets. In H. Gimple, N. R. Jennings, G. Kersten, A. Ockenfels, and C. Weinhardt, editors, *Negotiation and Market Engineering*, volume 2 of *LNBIP*. Springer, 2007.

J. Mäkiö and I. Weber. Modeling approach for auction based markets. In *The 2005 Symposium on Applications and the Internet*, pages 400–403, 2005.

I. ManyBrain. Mailinator. URL `http://sogetthis.com`.

R. Maranzato, M. Neubert, A. M. Pereira, and A. P. Lago. Feature extraction for fraud detection in electronic marketplaces. In *Proceedings of the 2009 Latin American Web Congress (la-web 2009)*, LA-WEB '09, pages 185–192, Washington, DC, USA, 2009. IEEE Computer Society. ISBN 978-0-7695-3856-3. doi: http://dx.doi.org/10.1109/LA-WEB.2009.20. URL `http://dx.doi.org/10.1109/LA-WEB.2009.20`.

R. Maranzato, A. Pereira, M. Neubert, and A. P. do Lago. Fraud detection in reputation systems in e-markets using logistic regression and stepwise optimization. *SIGAPP Appl. Comput. Rev.*, 11:14–26, June 2010. ISSN 1559-6915. doi: http://doi.acm.org/10.1145/1869687.1869689. URL `http://doi.acm.org/10.1145/1869687.1869689`.

R. D. McKelvey and T. Page. Public and private information: An experimental study of information pooling. *Econometrica*, 58(6):1321–1339, November 1990.

R. Molich and J. Nielsen. Improving a human-computer dialogue: What designers know about traditional interface design. *Communications of the ACM*, 33(3), March 1990.

A. E. Monge. Matching algorithms within a duplicate detection system. *IEEE Data Engineering Bulletin*, pages 1–7, 2000.

J. Moody, D. A. McFarland, and S. Bender-deMoll. Visualizing Network Dynamics. *American Journal of Sociology*, 110(4):1206–1241, January 2005. XXX.

NASD. Sonar (ads), 2008. `http://www.aaai.org/Pressroom/Releases/release-03-0917.php`.

F. . National White Collar Crime Center. Internet crime report. Technical report, Internet Crime Complaint Center, 2006.

Nestor. Prism efraud, 2008. `http://findarticles.com/p/articles/mi_m0EIN/is_2000_Feb_22/ai_59581365`.

D. Neumann. *Market Engineering - A Structured Design Process for Electronic Markets*. PhD thesis, Fakultät für Wirtschaftswissenschaften, Universität Karlsruhe (TH), Karlsruhe, 2004.

D. Neumann, J. Mäkiö, and C. Weinhardt. Came - a tool set for configuring electronic markets. In *ECIS*, 2005.

L. NFC Global. Nfc sentinel, 2008. http://www.nfcglobal.com/Sentinel.asp.

J. Nielsen and R. Molich. Heuristic evaluation of user interfaces. In *Proceedings ACM Conference Computer Human Interaction (CHI 1990)*, pages 249–256. ACM, April 1990.

U. C. Office. Identity fraud: A study. *Home Office*, 2002. URL http://www.homeoffice.gov.uk/docs/id_fraud-report.pdf.

R. Oprea, D. Porter, C. Hibert, R. Hanson, and D. Tila. Can manipulators mislead prediction market observers? November 2006. URL http://econ.ucsc.edu/faculty/roprea/manipEX2.pdf.

G. Ortner. Forecasting markets - an industrial application: Part i. Working paper, TU Vienna, Dep. of Managerial Economics and Industrial Organization, March 1998.

M. Ottaviani and P. N. Sørensen. Outcome manipulation in corporate prediction markets. *Journal of the European Economic Association*, page forthcomming, 2007. URL http://faculty.london.edu/mottaviani/omicpm.pdf.

D. M. Pennock. A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 170–179, New York, NY, USA, 2004. ACM. ISBN 1-58113-711-0. doi: http://doi.acm.org/10.1145/988772.988799.

D. M. Pennock, S. Lawrence, C. L. Giles, and F. Å. Nielsen. The power of play: Efficiency and forecast accuracy in web market games. Technical report, NEC Research Institute, Princeton, New Jersey, 2000. URL http://artificialmarkets.com/. A brief version appears in Science 291: 987-988, February 9, 2001.

C. Phua, V. Lee, K. Smith-Miles, and R. Gayler. A comprehensive survey of data mining-based fraud detection research. *Artificial Intelligence Review*, 2005.

C. R. Plott and K.-Y. Chen. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working Papers 1131, California Institute of Technology, Division of the Humanities and Social Sciences, Mar. 2002. URL http://ideas.repec.org/p/clt/sswopa/1131.html.

P. M. Polgreen, F. D. Nelson, and G. R. Neumann. Use of prediction markets to forecast infectious disease activity. *Clinical Infectious Diseases*, 44(2):272–279, 2007.

T. S. Porter, K. Schueller, T. Riley, R. Ruffer, and E. Usip. Marketsim: A simulated economy for microeconomics. *The Journal of Economic Education*, 37(4):483–483, Fall 2006. Heldref Publications.

C. Potts. Software-engineering research revisited. *Software, IEEE*, 10(5):19–28, September 1993. ISSN 0740-7459. doi: 10.1109/52.232392.

Progress Software Limited. Apama. online, 02 2008. URL http://www.apama.com/. Event Processing Platform.

H. C. Purchase, R. F. Cohen, and M. James. Validating graph drawing aesthetics. In F. J. Brandenburg, editor, *Graph Drawing 1995*, pages 435–446, Passau, Germany, 1996. Springer.

J. R. Quinlan. C4.5: Programs for machine learning. *Machine Learning*, 16(3):235–240, 1993. doi: 10.1007/BF00993309. URL http://www.springerlink.com/content/v986m1562062hk51.

ReD (Retail Decisions). Red shield, 2008. http://www.redplc.com/red209.asp.

P. W. Rhode and K. S. Strumpf. Historical presidential betting markets. *Journal of Economic Perspectives*, 18(2):127–142, Spring 2004.

P. W. Rhode and K. S. Strumpf. Manipulating political stock markets: A field experiment and a century of observational data. workingpaper, Januar 2007. URL http://www.unc.edu/~cigar/.

J. Robinson. *The Laundrymen: Inside Money Laundering, the World's Third Largest Business*. Pocket Books, London and New York, 1998.

D. Rolli, D. Neumann, and C. Weinhardt. A minimal market model in ephemeral markets. In M. Núñez, Z. Maamar, F. L. Pelayo, K. Pousttchi, and F. Rubio, editors, *FORTE Workshops*, volume 3236 of *Lecture Notes in Computer Science*, pages 86–100. Springer, 2004. ISBN 3-540-23169-2.

C. Rostoker. Interactive visualization of the market graph. Technical report, University of British Columbia, January 2006.

A. E. Roth. The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002. Fischer Schultz lecture.

S. F. Roth and J. Mattis. Data characterization for intelligent graphics presentation. In *Proceedings of the Conference on Human Factors in Computing Systems*, pages 193–200, April 1990.

F. Route. Validis, 2008. http://www.future-route.com/the_applications.

RuleQuest Research Pty Ltd. Gritbot. online, January 2008. URL http://www.rulequest.com/.

J. Schröder. *Manipulation in Prediction Markets - Analysis of Trading Behaviour Not Conforming With Trading Regulations*. Universität Karlsruhe (TH), 2009.

SearchSpace. Fraudoffice, 2008. http://telephonyonline.com/wireless/mag/wireless_detection_connection/.

Searchspace. Monitars, 2008. http://www.fortent.com/clients/index.php.

T. E. Senator. Ongoing management and application of discovered knowledge in a large regulatory organization: a case study of the use and impact of NASD regulation's Advanced Detection System (RADS). In *KDD '00: Proceedings of the sixth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 44–53, New York, NY, USA, 2000. ACM. ISBN 1-58113-233-6. doi: http://doi.acm.org/10.1145/347090.347102.

T. E. Senator, H. G. Goldberg, J. Wooton, M. A. Cottini, A. F. U. Khan, C. D. Klinger, W. M. Llamas, M. P. Marrone, and R. W. H. Wong. The Financial Crimes Enforcement Network AI System (FAIS) Identifying Potential Money Laundering from Reports of Large Cash Transactions. *AI Magazine*, 16(4):21–39, 1995.

T. E. Senator, H. G. Goldberg, P. Shyr, S. Bennett, S. Donoho, and C. Lovell. The NASD regulation advanced detection system: integrating data mining and visualization for break detection in the NASDAQ stock market. pages 363–371, 2002.

E. Servan-Schreiber, J. Wolfers, D. M. Pennock, and B. Galebach. Prediction markets: Does money matter? *Electronic Markets,*, 14(3):243–251, 2004.

H. S. Shah, N. R. Joshi, A. Sureka, and P. R. Wurman. Mining for bidding strategies on ebay. *Lecture Notes on Artificial Intelligence*, 2003.

G. Shmueli and W. Jank. Visualizing online auctions. *Journal of Computational and Graphical Statistics*, 14(2):299–319, 2005.

G. Shmueli, W. Jank, A. Aris, C. Plaisant, and B. Shneiderman. Exploring auction databases through interactive visualization. *Decision Support Systems*, 42(3):1521–1538, 2006.

M. Sipser. *Introduction to the Theory of Computation*, chapter 7.1 Measuring Complexity, pages 226–228. PWS Publishing, 1997.

SIRCA. Sirca report puts identity fraud at $1 billion. Technical report, Securities Industry Research Centre of Asia-Pacific Ltd., 2003.

Smarts Group International Pty Ltd. Smarts. online, January 2008. URL http://www.smartsgroup.com.

Smartsgroup. Smarts, 2008. http://smartsgroup.com/page.aspx.

M. A. Smith and B. Leigh. Virtual subjects: Using the internet as an alternative source of subjects and research environment. *Behavior Research Methods, Instruments, & Computers*, 29(4):496–505, 1997. URL http://www.psychonomic.org/search/view.cgi?id=127.

B. T. Solutions. Watchdog, 2008a. http://www.nofraud.com/Templates/Article10.aspx?PageID=d2039a53-a301-48f7-b4c0-9a942ab91eb1.

C. C. Solutions. Magnify pattern:detect, 2008b. http://claimsolutions.choicepoint.com/.

M. I. Solutions. Vector fraud solution, 2008c. http://www.metavanteimage.com/.

I. Sommerville, T. Rodden, P. Sawyer, R. Bentley, and M. Twidale. Integrating ethnography into the requirements engineering process. *Requirements Engineering, 1993., Proceedings of IEEE International Symposium on*, pages 165–173, 4-6 Jan 1993. doi: 10.1109/ISRE.1993.324821.

StatConsulting. Advancedminer, 2008. http://www.statconsulting.eu/telecommunication/fraud_detection.html.

B. Stefano and F. Gisella. Insurance Fraud Evaluation: A Fuzzy Expert System. In *Proc. of IFSC01*, 2001.

R. Stream. Rootstream detect, 2008. http://www.rootstream.com/products.aspx.

M. Strehlitz. Vertrauen und technik helfen gegen anarchie im enterprise 2.0. *Computer Zeitung*, (27):34, 29th, June 2009.

Subex. Nikira, 2008. http://www.subexworld.com/productsandsolutions/products/nikira.html.

K. Sugiyama. *Graph Drawing and Applications - For Software and Knowledge Engineers*, volume 11 of *Series on Software Engineering and Knowledge Engineering*. World Scientific, 2002.

N. Technologies. Minotaur, 2008a. http://www.neuralt.com/fraud_management.html.

N. Technologies. Neural technologies decider, 2008b. http://www.neuralt.com/credit_risk_management.html.

I. Tiburion. Linkexplorer, 2008. http://www.tiburoninc.com/solutions/link-analysis.asp.

Tildenwoods. Centrifuge, 2008. http://www.tildenwoods.com/.

D. Turo and W.-A. Jungmeister. Adapting treemaps to stock portfolio visualization. Technical report, University of Maryland (College Park, Md.), 1992.

G. Tziralis and I. Tatsiopoulos. Prediction markets: An extended literature review. *The Journal of Prediction Markets*, 1:75–91(17), 2007. URL http://www.ingentaconnect.com/content/ubpl/jpm/2007/00000001/00000001/art00006.

S. Universities. Data mining in identity crime prevention, 2008. http://www.bsys.monash.edu.au/people/cphua/.

J. Vila. Simple games of market manipulation. *Economic Letter*, 29(1):21–26, 1989.

C. von Altrock. Fraud detection. In *Fuzzy Logic and Neurofuzzy Applications in Business and Finance*, pages 286–294. Prentice Hall, New Jersey, 1997.

G. A. Wang, J. J. Xu, and H. Chen. Using social contextual information to match criminal identities. *hicss*, 4:81b, 2006. ISSN 1530-1605. doi: http://doi.ieeecomputersociety.org/10.1109/HICSS.2006.525.

Y. Wang and S. Madnick. The interdatabse instance identification problem in integrating autonomous systems. In *Proceedings in the 5th International Conference on Data Engineering*, Los Angeles, CA, 1989. IEEE Computer Society.

T. Washio and H. Motoda. State-of-the-art of Graph-based Data Mining. *SIGKDD Explorations*, 5(1):pp. 61–70, 2003.

S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, 1994.

M. Weatherford. Mining for fraud. *Intelligent Systems, IEEE*, 17(4):4–6, Jul-Aug 2002. doi: 10.1109/MIS.2002.1024744. URL `http://ieeexplore.ieee.org/iel5/5254/22033/01024744.pdf`.

C. Wehrend, S.; Lewis. A problem-oriented classification of visualization techniques. *Proceedings of the First IEEE Conference on Visualization*, pages 139–143, 469, 23-26 Oct 1990. doi: 10.1109/VISUAL.1990.146375.

C. Weinhardt, C. van Dinther, K. Kolitz, J. Mäkiö, and I. Weber. meet2trade: A generic electronic trading platform. In *Proceedings of the 4th Workshop on e-Business (WEB 2005)*, Las Vegas, USA, 2005.

C. Weinhardt, C. van Dinther, M. Grunenberg, K. Kolitz, M. Kunzelmann, J. Mäkiö, I. Weber, and H. Weltzien. *CAME-Toolsuite meet2trade - auf dem Weg zum Computer Aided Market Engineering*. Universitätsverlag Karlsruhe, Karlsruhe, 2006a.

C. Weinhardt, C. van Dinther, M. Grunenberg, K. Kolitz, J. Mäkiö, I. Weber, and H. Weltzien. *CAME-Toolsuite meet2trade - auf dem Weg zum Computer Aided Market Engineering*. University Press Karlsruhe, 2006b.

G. Weiss. Mining with Rarity: A Unifying Framework. *SIGKDD Explorations*, 6(1): pp. 7–19, 2004.

I. H. Witten and E. Frank. *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann, San Francisco, 2005.

WizSoft. Wizrule, 2008a. `http://www.wizsoft.com/default.asp?win=8&winsub=8`.

WizSoft. Wizwhy, 2008b. `http://www.wizsoft.com/default.asp?win=7&winsub=7`.

E. Wolfers, J..AND Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18 (2):107–126, 2004.

J. Wolfers and E. Zitzewitz. *Five Open Questions about Prediction Markets*. AEI-Brookings Joint Center for Regulatory Studies, 2006. URL `http://aei-brookings.org/admin/authorpdfs/page.php?id=1305`. Information Markets: A New Way of Making Decisions.

Xtract. Xtract autoscore, 2008. `http://www.xtract.fi/index.php/browse/index?category=1&subcategory=1&id=29`.