Expert Systems with Applications 39 (2012) 9966–9985

Contents lists available at SciVerse ScienceDirect

Expert Systems with Applications

journal homepage: www.elsevier.com/locate/eswa

The design of FFML: A rule-based policy modelling language for proactive fraud management in financial data streams

Michael E. Edge, Pedro R. Falcone Sampaio *

Manchester Business School, University of Manchester, Booth Street East, Manchester M15 6PB, United Kingdom

ARTICLE INFO

Keywords:
Fraud detection
Financial information systems
Rule-based expert systems
Information security
Domain specific policy-based languages

ABSTRACT

Developing fraud management policies and fraud detection systems is a vital capability for financial institutions towards minimising the effect of fraud upon customer service delivery, bottom line financial losses and the adverse impact on the organisation's brand image reputation. Rapidly changing attacks in real-time financial service platforms continue to demonstrate fraudster's ability to actively re-engineer their methods in response to ad hoc security protocol deployments, and highlights the distinct gap between the speed of transaction execution within streaming financial data and corresponding fraud technology frameworks that safeguard the platform. This paper presents the design of FFML, a rule-based policy modelling language and encompassing architecture for facilitating the conceptual level expression and implementation of proactive fraud controls within multi-channel financial service platforms. It is demonstrated how a domain specific language can be used to abstract the financial platform into a data stream based information model to reduce policy modelling complexity and deployment latencies through an innovative policy mapping language usable by both expert and non-expert users. FFML is part of a comprehensive suite of assistive tools and knowledge-based systems developed to support fraud analysts' daily work of designing new high level fraud management policies, mapping into executable code of the underpinning application programming interface and deployment of active monitoring and compliance functionality within the financial platform.

1. Introduction

Fraud can be defined as "wrongful or criminal deception intended to result in financial or personal gain" (Oxford, 2008). The potential for large scale monetary gain coupled with the rapidly changing nature of financial information service platforms presents a unique range of opportunities for the diligent fraudster. Detection of fraudulent behaviour within large data sets has become an extensive area of research within the last decade resulting in an extensive body of academic literature (Kou, Lu, Sirwongwattana, & Huang, 2004) and fraud solution vendors (Experian, 2011; FICO, 2011; RSA, 2011) towards minimising the effects of fraud within the financial service marketplace. Yet rates of fraud within the financial domain continue to escalate with recent fraud figures illustrating persistent high levels, despite the deployment of pioneering fraud prevention schemes and technologies. The latest Annual Fraud Indicator report from the National Fraud Authority in the United Kingdom (NFA, 2011) has put the loss to the UK economy from all types of fraud at £38.4 billion (US$ 61 billion, approximately), around £8 billion more than the 2010 estimate. Also according to the NFA report, latest estimates on online banking fraud in the United Kingdom increased by 14%, from £53 million in 2008 to £60 million in 2009. A recent study of US card payment fraud by Sullivan (2010) also indicates that the costs linked to online payment fraud (lost sales, direct payment fraud losses, and fraud management) rose from 2000 to 2009, with 2009 costs estimated at US$ 3.3 billion (1.2% of sales revenue). Sullivan (2010) also indicates that "the fraud loss rate for the US appears to be higher than that of Australia, France, Spain and the UK".

Automated fraud detection research has its foundations within the knowledge discovery domain of artificial intelligence, providing a wealth of literature over a vast number of application domains. Early contributions within insurance (Viaene, Dedene, & Derrig, 2005), telecommunications (Boukerche & Notare, 2000; Burge et al., 1997) and credit card fraud detection (Aleskerov, Freisleben, & Rao, 1997; Brause, Langsdorf, & Hepp, 1999; Ghosh & Reilly, 1994; Maes, Tuyls, Vanschoenwinkel, & Manderick, 2002) demonstrate the capability of supervised data mining techniques for extraction of known fraud scenarios based upon previously experienced and labelled fraud cases. Subsequent research efforts have converged towards reducing the reliance of employed algorithms upon labelled training data to complement the increasing speed of underpinning financial information systems and associated temporal window constraints. Phua, Lee, Smith, and

* Corresponding author. Tel.: +44 0 1613063349.
E-mail addresses: michael@edgem.net (M.E. Edge), Pedro.Sampaio@manchester.ac.uk (P.R. Falcone Sampaio).

0957-4174/$ - see front matter © 2012 Elsevier Ltd. All rights reserved.
http://dx.doi.org/10.1016/j.eswa.2012.01.143

Gayler (2005) describe how "if the incoming data stream has to be processed immediately, in an event driven system, or labels are not readily available, then semi-supervised or unsupervised approaches are the only data mining options".

Technology adoption has continued to drive evolution of financial information systems into multi-channel service infrastructures supporting global user delivery of financial products and services. Extensive research effort has therefore been invested in parallel advancement of fraud technologies to encompass emerging service infrastructures and the reduced time windows within which large scale financial fraud can now be performed. Demands associated with event based data evaluation and the increasing costs of fraudulent cases however have rapidly exceeded the capabilities of traditional knowledge discovery models, stimulating research within the development of real-time data processing technologies to support the fraud detection operation. STREAM (Arasu et al., 2003), Aurora (Abadi et al., 2003), Borealis (Abadi et al., 2005) and Telegraph (Chandrasekaran et al., 2003) all demonstrate the benefits that can be achieved through direct evaluation of information upon incoming event streams prior to data storage. Luckham (2002) in particular demonstrates the ability to monitor and trigger preventive actions in response to system events within large scale information system architectures. StreamBase (2011), Sybase (2011) and ruleCore (2011) also represent commercial developments in response to the growing demand for real-time business analytics.

Massey (2005) describes traditional fraud detection systems as "more reactive than proactive". In *reactive fraud management* knowledge discovery techniques based upon data mining (Phua et al., 2005) are implemented to perform algorithmic processing and data analysis over static data repositories. Fraudulent instances are identified using pre-defined fraud libraries or as anomalous behaviour against the accounts behavioural history. Implementation of a 'store now, query later' approach however significantly increases the incurred fraud detection latency due to the requirement of transactional data within the assessed data store prior to application of data analysis techniques. Triggering of a preventive response may therefore only be undertaken following transaction completion and movement of the associated monetary value.

In *proactive fraud management* newly arriving requests and click streams are analysed "on-the-fly" prior to transaction completion (Abadi et al., 2003; Arasu et al., 2003), enabling the identification of ambiguous instances prior to the movement of any financial value. Preventive actions can therefore be triggered in questionable instances to further verify the identity of the initiating user (Bicakci & Baykal, 2003), or declining of requests in high risk transaction scenarios. Fraud management consequently becomes an integrated aspect of the implemented authorisation process rather than a post transactional analysis operation over static data records.

Fraud statistics however continue to rise despite extensive development towards empowering financial institutions with innovative information architectures and technologies. While extensive literature exists within the system level design of both academic and commercial fraud technologies, a distinct gap exists regarding their functional deployment and enforcement as integral components of an organisation's fraud management operation. The absence of application level tools for supporting emerging data processing models is a major concern given that an organisation's fraud strategy approach and active policy repository must be continually revised in response to prevailing fraud patterns and behaviours. Furthermore, a current absence of standards for fraud policy definition precludes an institution's ability to share fraud policy data with analogous sector organisations towards reducing the latency associated with fraud threat discovery and deployment of associated fraud controls. Standards for fraud policy definition and dissemination therefore remain an area of significant challenge towards facilitating effective detection and regulation of fraud within the financial domain and associated fraud management operations.

It is possible to categorise the main challenges associated with the deployment of proactive fraud management within financial information systems as follows:

1. *Effective policy discovery*: Internal strategy development is a well established operational process with extensive academic literature. (Phua et al., 2005) and commercial tools (SAS, 2011; SPSS/IBM, 2011) to support extraction of knowledge from large transactional data sets. Yet financial institutions continue to incur extensive fraud losses due to the latency period experienced between deployment of new fraud control policies and identification through employed data discovery techniques. Furthermore, little work exists regarding the active circulation of emerging fraud threats prior to attack exposure and incurring the associated operational and financial costs. Significant work still remains in facilitating the active sharing of fraud knowledge between financial institutions to preclude emerging fraud attacks, and fraudsters' attempts to horizontally permeate their methods throughout analogous sector organisations.
2. *High level policy specification*: Challenges remain in the implementation of effective fraud controls using low level application programming interfaces associated with emerging data processing technologies. Timely construction, testing and deployment of required fraud controls are essential to capitalise upon the real-time evaluation of incoming data requests. Absence of adequate policy modelling tools significantly impedes an organisation's ability to preclude shifting fraudster behaviours through rapid fraud strategy realignment and deployment of optimal fraud controls.
3. *Flexible architectures for fraud detection and prevention*: Commercial solutions within the fraud detection space continue to supply approaches tailored for fraud detection within a single service channel or financial delivery platform (Entrust, 2011; FICO, 2011; RSA, 2011). Institutions hosting multiple service delivery channels therefore require multiple fraud technology solutions to achieve complete system wide coverage of the underlying financial service framework. The result is a highly fragmented fraud management architecture with a propensity to impede operational efficiency and cross-channel evaluation as business analysts must deploy fraud operations through numerous disparate application programming interfaces.
4. *Policy expression standards for fraud management*: Despite extensive work within the fraud domain, standards to support the common expression and enforcement of fraud policy controls remain an open research challenge. Bespoke application programming interfaces continue to be the de facto standard for researchers and practitioners in the deployment of required system functionality. Standards for fraud policy control are seen as analogous to database query languages such as SQL and OQL, with the potential for supporting transferable industry knowledge and analyst skills within the fraud management domain.

Technology innovation within financial service computing clearly dictates the need for real-time fraud analytics to complement the ever reducing time window within which large scale financial fraud can be performed. Effective computational tools and technologies therefore remain key components towards leveraging a proactive approach to fraud management and regulating fraud within rapidly evolving financial service platforms.

To address some of the challenges identified above, this paper presents a rule-based policy modelling language and encompassing architecture for facilitating the conceptual level expression