

Immediate Detection of Anomalies in Call Data - An Adaptive Intelligence Approach

Isaac O. Osunmakinde and Anet Potgieter
Agents Research Group
Department of Computer Science, Faculty of Science
University of Cape Town, Cape Town, Republic of South Africa
E-mail: (segun@cs.uct.ac.za and anet@cs.uct.ac.za)

Abstract - Anomaly detection in telecommunications call data tries to discover deviant behaviour of individual subscribers. Malicious behaviour has been noted as one of the key causes of such anomalies, which have consequently led to unquantifiable losses of revenue to many telecommunication networks world-wide. Although the intentions of most subscribers to these networks are unknown when making phone calls, their behaviour pattern is reflected in their call data. Recent studies have investigated the challenges of anomaly detection but have not given conclusive solutions to address this problem. To maximize the true positive rates and minimize false detection at an acceptable level, this paper presents the Telecommunications Anomaly Detection System (TADS) which efficiently facilitates immediate detection that will help call analysts and managers with adaptive decision-making. We developed and implemented the TADS which uses Behavioural Bayesian Networks (BBNs) to address this challenge. The empirical evaluation results obtained using real world land-line call data for subscribers of a local Telecommunication Services Provider (TSP) demonstrated that the TADS can detect and act upon anomalies in telecommunication networks as it happens.

Keywords: Anomaly Detection, Bayesian Network, Subscriber / Network management.

I. INTRODUCTION

Since the telephone is one of the fastest means of communication, many subscribers use contract and pre-paid phones for crucial business activities. Security deficiency is observed in the fixed and mobile networks. In 2001, half a million fixed lines in South Africa were disconnected, many as a result of fraudulent activities [9]. The security measures that include the use of a Mobile Identification Number (MIN) / Equipment Serial Number (ESN) and a Personal Identification Number (PIN) were meant to prevent spoofers (or unauthorised users) from gaining access into a subscriber's phone account. These measures are deficient due to unlawful methods [13], such as eavesdropping, cracking, shoulder surfing, and social engineering [14], [15]. To worsen the situation, the MIN / ESN numbers for digital phones (e.g. GSM) that was initially well encrypted during transmission, was later cracked (made simpler) because of the decoding disparity from one country network to another [15]. As a result, internetwork linking countries have lax security which is easy to breach. This insecurity allows the spoofers or hackers to make anomalous calls at the expense of the lawful phone owners.

The term anomaly refers to an outlier and to suspicious data that can easily be spotted in small datasets but is hidden and requires intelligent detection in the

massive amounts of call data in telecommunications environments. Common inconsistencies in call data are caused by customer churn or attrition, failure in networks, potential fraud, deliberate or unintended expensive mistakes made by telecommunications' workers, bad debt risk, or even improper disconnection of communication lines [2], [3], [4]. Malicious behaviour has been noted as one of the key causes for such anomalies, which have consequently led to unquantifiable loss of revenue to many telecommunication networks world-wide [9], [17]. The most common cause of such malicious behaviour is potential fraud, because it is committed intentionally and it is difficult to combat. To complicate the situation further, the network carriers often do not want to admit that fraud exists as an anomaly in their systems, so that their subscribers do not suspect that fraud is a significant problem. If they do acknowledge it as a significant problem, it might cause churn or cause more subscribers to try to commit fraud. Instead, they prefer to solve this problem intelligently. If the subscribers suspect fraud but nonetheless keep paying for debts that they did not incur and for services that they did not receive, and if the Telecommunication Service Provider (TSP) does not find a solution to the problem, these customers may decide to seek alternative competing service providers. Also, service providers are greatly challenged when subscribers, who feel uncomfortable with their services including for example; over-charging, not answering queries promptly, or generally bad customer service, change to competing carriers. These changes contribute to suspicious and sudden change in call patterns that is reflected in their call data. It results in the loss of revenue, which is mostly attributed to anomalies. This could quickly put a telecommunication company out of business. For these reasons, it is obvious that anomaly prevention is defeated. This necessitates the use of our proposed immediate detection techniques for existing subscribers who experience anomalies.

A reduced lag-time (difference between time of call and detection time) will enable service providers to detect and act upon anomalies faster, which will reduce the damage caused by these anomalous behaviours without the appropriate intervention. Also, an anomaly detection method can work well for a finite number of known anomalies but due to the difference in behaviours of subscribers as a result of uncertainty in the use of phone services, a system to detect infinitely number of unknown anomalies is preferable. Accomplishing this objective is very challenging. Therefore, a TSP that can detect the trustworthiness of any call record immediately after a call session is ended will flourish and create a level of confidence between itself and its subscribers. On the basis of our previous work on individual BBNs [10], we

developed and implemented an effective TADS that detects any anomalies and minimises false alarms to an acceptable level. It comprises the following four components: Call Sampling Probability, Call Prediction, Call Detection, and Degrees of Detection. Distinguishing between calls made to new destination numbers and detection of anomalous calls is difficult but the TADS uses anomaly indicators [10] to address this problem effectively. In principle, the Call Detail Records (CDR) that describe subscribers' user profiles generally contain relevant attributes, such as a caller's number, location, duration, destination number, date and time of call. Assignment of equal weights to all these attributes, even though some carry more information than others, is an approximate assumption. Unequal weighting is determined as indicators, which improve the appropriateness of TADS and the telecommunication revenue losses will be significantly reduced. To realise this constructive idea and to facilitate a trusted relationship, a warning notification through alternative phone, sms, or e-mail can be sent by the TSP to subscribers whose account reflects abnormal call patterns.

Acronyms used in this paper were defined at first occurrences. In the remaining part of the paper, section two describes the probabilistic network and behavioural modelling, section three describes the TADS, section four presents the implementation and evaluation results, section five reviews the related work and section six describes the concluding remarks and future work.

II. PROBABILISTIC NETWORK AND BEHAVIOURAL MODELLING

The probabilistic networks described in this section are Bayesian networks (BNs). This section also describes how we use Genetic Algorithms (GA) for individual behavioural modelling.

A. Bayesian Networks and Reasoning

A BN model is a multivariate probability [11] distribution that is used to define qualitative and quantitative relationships between random variables X_1, \dots, X_n . The relationships describe the causalities or dependencies of variables. The qualitative knowledge is captured in the connection of nodes (variables) and the directed arcs as shown in Figure 2.1 while the quantitative knowledge is captured in the Conditional Probability Tables (CPTs). Learning CPTs is described by Murphy [19] and Russel [20]. A Bayesian belief network formally requires discrete random values such that if there exists random variables X_1, \dots, X_n with each having a set of values x_1, \dots, x_n then, their joint probability density distribution is defined in equation 2.1;

$$\Pr(X_1, \dots, X_n) = \prod_{i=0}^n \Pr(X_i | \pi(X_i)) \quad 2.1$$

where $\pi(X_i)$ represents a set of probabilistic parent(s) of child X_i . [11]. A parent variable refers to the cause while a child variable means the effect. Figure 2.1 shows the Directed Acyclic Graph (DAG) representation of a Bayesian network model, BBN, which describes the behaviour of a telecommunication subscriber used to detect anomalies in phone calls [10]. The BBN shows mined relationships between destination numbers and call duration attributes (nodes), relationships between destination networks and location attributes. and is used for anomaly detection in the

presence of uncertainty about subscriber behaviour. In this paper, we want to reason with expressions 2.2 and 2.3.

$$\Pr(a \text{ call} = \text{regular} | \text{Subscriber's behaviour}) \quad 2.2$$

$$\Pr(a \text{ call} = \text{anomalous} | \text{Subscriber's behaviour}) \quad 2.3$$

Expression 2.2 represents the conditional probability that a call is regular, given a subscriber's behaviour and expression 2.3 is the conditional probability that a call is anomalous, given same subscriber's behaviour. An important reasoning functionality used is Bayesian inference [22] which adds to the power of our TADS.




Figure 2.1: Subscriber 145521137 Call Behaviour.




Figure 2.2: Subscriber 145521198 Figure 2.3: Subscriber 145571179




Figure 2.4: Subscriber 145571042 Figure 2.5: Subscriber 145571030

B. Individual Modelling With Genetic Algorithms

Intuitively, learning a BN from telecommunications data can be referred to as discovering the behavioural call patterns (models) of subscribers. Figures 2.1-2.5 show that subscribers' behaviours are reflected in their call data.

A GA which learns the individual subscriber's behaviour from call data has been presented by Osunmakinde and Potgieter [10]. The algorithm uses information-theoretic measures (e.g. Minimum Description Length) and mathematical components (e.g. PowerSet in set theory) as genetic operators and as a means of balancing between efficiency and decomposability. As a result of this, the causalities or dependencies discovered in subscribers' models are suitable for detecting anomalies using the associated anomaly indicators. In [10], we formally described the identification of anomaly indicators from the most interesting / significant call attributes of models as primary indicators. It varies from one subscriber to another as one subscriber's regular call may be another subscriber's anomalous call. This technique makes it possible for our TADS to differentiate between a call made to a new destination number and an anomalous call. In addition to the individual models presented in [10], Figures 2.1 to 2.5 are more examples of individual subscribers' models learnt from call data using the GA.

III. TELECOMMUNICATIONS ANOMALY DETECTION SYSTEM (TADS)

The TADS is a differential model, which includes sampling and inferential tests to detect anomalies in call data. Differential analysis adapts to dynamic behavioural call patterns associated with different telephone subscribers. For example, a behavioural pattern detected as an anomalous call for subscriber X may be accepted as a regular call for subscriber Y. This approach is more efficient than a fixed trigger system that raises an alert when any differences are suspected in call patterns. Our TADS is based on the interesting recognition of an event discussed by Jaroszewick [1]. For qualitative detection, we improve the idea to formalise this new detection approach. The next subsection describes the system model.

A. The Descriptive Approach of TADS

The system model illustrated in Figure 3.1 gives an overview of the TADS. It uses Bayesian learning, Bayesian inference and differential analysis [11], [20]. During Bayesian learning, the TADS uses the historical call dataset (HCD) to train an evolved individual BBN from the GA and the quantitative behaviour is captured as conditional probability tables (CPTs). The TADS acts on a trained BBN to make a decision about a current call record immediately after a call session ends. Call data identified as regular is added to the HCD to re-train (adapt) the BBN model, while the anomalous call alerts the call analysts to investigate. The anomalous calls can be confirmed by the TSP as regular or anomalous. This could be assisted with interrogation through e-mail or the subscriber's alternative phones. If it is confirmed as regular, it will be included into the HCD for training. If it is confirmed as anomalous, then it is totally excluded from the training set.

The addition of regular calls to the HCD is used to improve the accuracy of the BBN before the TADS acts upon it. The use of this re-training is strongly encouraged to minimise error rates (false alarms).

B. Components of TADS

The components that serve as the building blocks of TADS subsystem in Figure 3.1 are the following:

Call Sampling Probability

Call sampling probability is the probability of $E = \bar{e}$ sampled from the training set. In practice, \bar{e} refers to a call instance that is sampled from the training set. Sampled probability is represented as $\Pr_E^{\Xi}(\bar{e})$ and can be computed as:

$$\Pr_E^{\Xi}(\bar{e}) \Rightarrow \Pr(E = \bar{e}) = \frac{\#(\text{instances } \bar{e})}{\#(E)} \quad 3.1$$

where # means number.

Equation 3.1 is differential as long as call patterns of subscriber's behaviour change and it is therefore used as a form of call threshold to guide the degrees of detection. In the implementation of our TADS, the BBN is trained by computing its CPT using the training set.

Call Prediction

Call prediction represented as $\Pr_E^{BN}(\bar{e})$ is the probability

of $E = (\bar{e})$ inferred from the Bayesian network. This is computed using the Bayesian inference engine of JavaBayes [22] embedded in our TADS implementation.

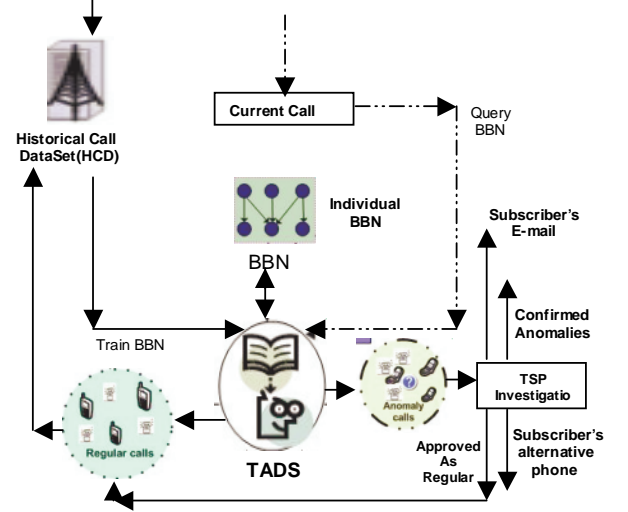


Figure 3.1: A System Model for the TADS.

Call Detection

Call detection is the decision making aspect of our TADS using the computed sampled and inferred probabilities. As a scientific test, this module suspects that a phone call record is either *regular* or *anomalous*. We therefore define call detection as follows:

$$\Pr_E^{BN}(\bar{e}) \geq \Pr_E^{\Xi}(\bar{e}) \Rightarrow \text{Regular Call Suspected.}$$

That is, if the inferred probability is greater or equal to the sampled probability, then a regular call is suspected. If the inferred probability is less than the sampled probability, then, an anomalous is suspected. That is:

$$\Pr_E^{BN}(\bar{e}) < \Pr_E^{\Xi}(\bar{e}) \Rightarrow \text{Anomalous Call Suspected.}$$

Automatic detection of anomalous calls can reduce the workload of telecommunications service providers. *Regular* implies that it follows the historical call patterns behaviour of the subscriber, while *anomalous* means it does not. The next module describes the level of confidence of call detection.

Degrees of Detection

Degrees of detection refer to the level of belief with regard to the call detection decisions. This result is expressed as a percentage and it will help call analysts and managers to make final decisions. A deviation *below* the sample probability is computed using equation 3.2. A deviation *above* the sample probability is computed using equation 3.3. The elements of equations 3.2 and 3.3 use differential analysis due to their regular update in input values. The results imply that we are **X.xx%** sure that a call is *anomalous* or *regular*. The implementation results of the TADS are shown in section four.

IV. IMPLEMENTATION RESULTS & EVALUATIONS

This section shows the immediate call detection results for individual subscribers and the performance evaluations of the TADS. We obtained real world land-line call data for over 160 TSP subscribers, and more than 73,000 calls with 9 attributes. A few representative results are presented below.

By masking the subscribers' phone numbers, we have ensured the confidentiality of network carriers based on the policies to protect telecommunication customers.

$$\left(\frac{\Pr_E^{\Xi}(\bar{e}) - \Pr_E^{BN}(\bar{e})}{\Pr_E^{\Xi}(\bar{e})} \times 100 \right) \% = \mathbf{X.xx\%} \quad 3.2$$

$$\left(\frac{\Pr_E^{BN}(\bar{e}) - \Pr_E^{\Xi}(\bar{e})}{\Pr_E^{BN}(\bar{e})} \times 100 \right) \% = \mathbf{X.xx\%} \quad 3.3$$

A. Immediate Call Detection Results

The following are examples of the immediate call detection results from our implementations.

TADS Results for a Subscriber

Sample call data for subscriber 145521137 is shown in table 4.1. Our TADS trained and used the BBN model in Figure 2.1, to act upon the current call records. The anomalous and regular calls detected immediately are shown in Table 4.2. The table shows the call detection results for the phone calls that were made on *Saturday*. A detection took 0.42 Secs. All regular call events were used to re-train the BBN to update its intelligence. The last call event result that was detected as *regular* only has a degree of belief of 4.441%, which is due to the use of the secondary anomaly indicator rather than the primary anomaly indicator. This can possibly be an example of a call made to a new phone number because this particular call record has a characteristic behaviour reflected in the historical behaviour of this subscriber. Similar detection results were obtained for other individual subscribers, of which a few will be presented in this paper.

B. Performance Evaluations of The TADS

TSPs cannot easily make call data available to researchers; a scientific alternative used is the 90% & 10% cross validation techniques [20]. Two evaluation cases are presented below. We evaluated the accuracy of the TADS by comparing the *expected call data* and the *actual calls detected* for an individual subscriber.

TEST CASE ONE: ANOMALY DETECTION WITHIN A SUBSCRIBER'S PROFILE

In test case one, regular and anomalous calls are detected from every subscriber's call data. A training set consisting of 90% of a subscriber's call data was used while 10% was chosen at random as a test set. This process was repeated for five subscribers chosen at random and the TADS accuracies were computed and are shown as a confusion matrix below.

TADS Accuracies for a Subscriber

Using the 10% test set, the expected call data is shown in Table 4.3. This is because domain experts did not expect any anomalies. That is, none of the training records was known as anomalous before modelling the BBN. The confusion matrix generated from the implementations, for subscriber 145521137, is shown in Table 4.4. The TADS accuracy for the subscriber is 87.5%. This is the true positive rate while the false alarm rate is 12.5%. That is, the anomalous calls detected in this context are regarded as false call detections. Similar results were obtained for other subscribers and the accuracy results for subscribers 145521198, 145571179, 145571042 and 145571030 are

58.065% & 41.935%, 77.358% & 22.642%, 97.633% & 2.367%, and 93.846% & 6.154% respectively. The overall accuracy from the implementations is 82.880%. The interpretations for the degrees of beliefs of the TADS results are visualised in Figure 4.1. From the detection results, one can see that the true positive rates are maximised while the false alarms are minimised at an acceptable level.

Table 4.1: Observed training set for subscriber 145521137

x1	x2	x3	x4	x5	x6	x7
Wed	218-261	828050678	Ce	Voda-com	6.293-7.552	P
Fri	44-87	1023	Ss	0-50Km	0.0-1.258	P
Sat	131-174	835575170	Ce	MTN	3.776-5.034	P
Tue	0-44	145924180	Na	0-50Km	0.0-1.258	X
Mon	87-131	829778104	Ce	Voda-com	1.258-2.517	P
..

Legend: Ce = Cellular, Na = National, x1 = Call-date, x2 = Duration in seconds, x5 = Location, x6 = Call-cost, x4 = Destination-net, x3 = Destination-no, x7 = Peak (P)/off-peak (X), Ss = Special Services.

Table 4.2: TADS detects current calls for subscriber 145521137

Immediate Call Events	
[Saturday, 44-87, 828581169, Cellular, (Vodacom), 1.258-2.517, P]	Call Detection: Regular Call Suspected Degree of Belief: 83.489%
[Saturday, 0-44, 836431694, Cellular, (MTN), 1.258-2.517, P]	Call Detection: Anomaly Call Suspected Degree of Belief: 30.372%
[Saturday, 0-44, 834143703, Cellular, (MTN), 1.258-2.517, P]	Call Detection: Anomaly Call Suspected Degree of Belief: 65.186%
[Saturday, 87-131, 834105953, Cellular, (MTN), 1.258-2.517, P]	Call Detection: Regular Call Suspected Degree of Belief: 4.441%
..	..

Table 4.3: Expected call detection results for subscriber 145521137.

Expected	Total
Regular Calls	40
Anomaly Calls	0

Table 4.4: Confusion matrix results for subscriber 145521137

Expected / Predicted	Regular	Anomaly
Regular	35.0	5.0
Anomaly	0.0	0.0

TEST CASE TWO: DETECTION USING A TEST SET THAT INCLUDES INTRODUCED ANOMALIES

Similarly to test case one, test case two uses 90% of the HCD as a training set to model the individual BBN and 10% of anomalous test data (noise) was randomly generated. The noise records were generated as known anomalies using the properties of the real dataset and it was repeated for other subscribers. Since all the test records are known anomalies, 100% detected anomalies are expected. In this case, the calls detected as anomalies are regarded as true (correct) detection, while the calls detected as regular calls are false (incorrect) alarms. The accuracies of the detection are summarised in Table 4.5. In this case, the average accuracy of correctly detected anomalies is 74.754%, while incorrect

(error) detection is 25.246%. One can see in Table 4.5 that, for every subscriber, the true positive rate is greater than the error rate. We cannot have a group model for subscribers who behave in a similar way. If we do not have individual models to separate subscribers' behaviours so that detection is immediate, the false detection rates may be greater than the true positive rates. Figure 4.2 visualises these detection results. Table 4.6 summarises the two test cases for the overall average detections. One can see that our methodology is good as a proof of concept since the overall average detection accuracy is 78.817 %.

Table 4.5: The anomaly detection accuracies for eight subscribers

Call Number	True Positive Rate%	False Alarms%
145521137	70.0	30.0
145571179	88.889	11.111
145571030	75.0	25.0
145521198	66.667	33.333
145571046	72.727	27.273
145571050	77.778	22.222
145571055	63.634	36.366
145571042	83.333	16.667
Average accuracies	74.754 %	25.246 %

Table 4.6: Summary of results' accuracies for the two test cases

Test Cases	Accuracy (%)
Test Case One	82.880
Test Case Two	74.754
Overall Average Accuracy	78.817 %

V. RELATED WORK

A number of methods used for anomaly detection include: rule-based approaches, statistical techniques, neural networks, distance-based approaches, and Bayesian networks. It was presented in recent studies that unsupervised learning of Bayesian networks for anomaly detection, which was investigated in this research, is better than most detection methods [7]. This is because it does not need to preset types of fraudulent data and can detect new anomalies from call data.

The *rule-based* techniques work best with user profiles containing explicit information, where anomaly criteria are encoded as rules [4]. This technique is difficult to manage because it requires explicit rules which are labour-intensive and time consuming, and involves programming for every imaginable possible anomaly. The *statistical* methods of anomaly detection work well with univariate distributions, where the average call duration, the longest call duration and the average number of calls per day are usually compared with a pre-determined threshold [6]. If the number of calls for a day exceeds its normal behaviour (threshold), it is considered as anomalies which may not be true. The popular *neural network* technology [5] employs feature extraction, where summarised statistics are usually used to train models with Call Detail Records (CDR). The extracted features from CDRs are pre-classified as either a regular or an anomalous call. This technique can detect known anomalous calls in the training data. Hollmen [3] presented the use of a *Bayesian network* in anomaly detection by profiling users' behaviours and used known anomalous scenarios but did not address new types of anomalous calls. Also, the *Distance-based* approach is described in [2], [21]. In this method, an absolute distance

between a dataset and a defined point as a distance outlier is used. This approach focuses primarily on continuous datasets. Our approach differs from theirs as we can also process mixed datasets (see Table 4.1). This paper describes how individual models are learnt from data which are then used to detect anomalies immediately after a call session ends. This reduces the time lag with degrees of detection. Also, the differential threshold makes our new TADS adaptive and minimises false alarms to an acceptable level.

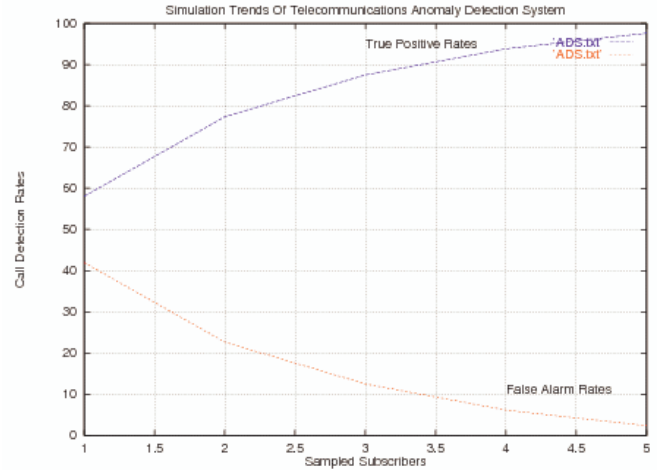


Figure 4.1: Prediction rates of TADS for test case one

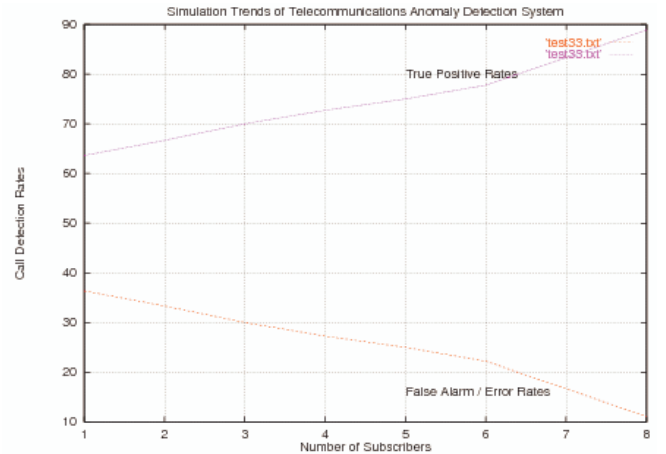


Figure 4.2: Prediction rates of TADS for test case two.

VI. CONCLUDING REMARKS

The most important contribution of this paper was the adaptive intelligence we achieved with the development of the TADS, having an average detection accuracy of 78.817%. The intelligence of the TADS is a result of the following: the use of individual BBNs, re-training of the BBNs, the use of a differential threshold, the use of anomaly indicators, and immediate action on current calls with unknown anomalies.

In our detection results above, the TADS aimed and was able to maximise true positive rates and minimise false alarms. The proven performance of our approach makes it suitable to applications in various areas such as [10]. We intend to apply it to other interesting application areas such as internet banking, business pattern recognition, stock & image analysis etc. We have established a proof of concept in telecommunications but currently working on its scalability. Our system can be deployed on networked workstations. We did not have access to mobile call data but

we are fortunate that our landline data contains relevant attributes that are equivalent to the mobile data. This paper is a summary of our work described in [8].

ACKNOWLEDGEMENT

Special thanks to Complex Adaptive Systems (Pty) for funding this research. Many thanks to the developers of the JavaBayes, and Gnuplot.

REFERENCES

- [1] Jaroszewick, S. & Simovici, D. (2004). Interestingness of Frequent Itemsets using Bayesian Networks as Background Knowledge. *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Pages: 178 - 186, ACM Press.
- [2] Bay, D. & Schwabacher, M. (2003). Mining distance-based outliers in near linear time with randomization and a simple pruning rule. In *proc. of 9th annual ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Pages 29-38.
- [3] Hollmen, J., Tresp, V., Taniguchi, M., & Haft, M. (1998). Fraud Detection In Communications Networks Using Neural And Probabilistic Methods. *Proceedings of the 1998 IEEE Int. Conf. in Acoustics, Speech and Signal Processing (ICASSP'98)*, Vol. 2, Pages 1241-1244.
- [4] Fawcett, T. & Provost, F. (1997). Adaptive Fraud Detection. *Data Mining and Knowledge Discovery*, Kluwer Academic Publishers, Boston, CA. Vol 1, Pages 291-316.
- [5] Bounsaythip, C. & Rinta-Runsala, E. (2001). Overview of Data Mining for Customer Behaviour Modelling. *Technical Report*, TTE1-2001-18, VTT Information Technology, Finland.
- [6] Michael, H., Lambert, D., Pinheiro, C., & Sun, D. (2002). Detecting Fraud in the Real World. *Handbook of massive data sets*, Pages 911-929, Kluwer Academic Publishers.
- [7] Kou, Y., Lu, C. Sirwongwattana, S. & Yo-Ping, H. (2004). Survey of Fraud Detection Techniques. *Networking, Sensing and Control, IEEE International Conference*, Page(s) 749-754 Vol.2, IEEE.
- [8] Osunmakinde, I. O. (2006). Intelligent Detection of Anomalies in Telecommunications Customer Behaviour. *M.Sc. thesis*, University of Cape Town, Computer Science Department.
- [9] Frayne, M. (2006). Interconnect billing- Making Telecommunications Work for Africa, *Annual Telecommunications Report*, Intec Telecom Systems. http://www.connect-world.com/Articles/old_articles/MikeFrayne.htm.
- [10] Osunmakinde, I. O. & Potgieter, A. (2006). Agent-Based Behavioural Modelling for Anomaly Detection in Call Data from Telecommunication Networks. *Proceedings of Southern African Telecommunications Networks and Applications Conference (SATNAC)*, Pages 8-9.
- [11] Pearl, J. (1988). Probabilistic reasoning in intelligent systems. *Networks of Plausible Inference*, Morgan Kaufmann Publishers.
- [12] Potgieter, A., April, K.A., Cooke, R.J.E. & Osunmakinde, I. O. (2006). Understanding Social Complexity, *Journal of Emergence: Complexity and Organization* (E:CO).
- [13] Brooks, T. & Davis, M. (1994). Are Your Phone Bills Fraud Free? *Security Management*, vol. 38, no. 4, Pages 67-68.
- [14] Delaney, P. (1993). Investigating Telecommunications Fraud. *Criminal and Civil Investigation Handbook*, 2nd ed., McGraw-Hill Inc. New York.
- [15] Gillian, D. (1999). The changing face of fraud : phone fraud, *3rd National Outlook Symposium on Crime in Australia: Mapping the Boundaries of Australia's Criminal Justice System* Rydges Hotel, Canberra.
- [16] Sequeira K. & Zaki M. (2002) Anomaly-based data mining for intrusions. In *ACM SIGKDD 02*.
- [17] Rosset, S., Murad, U., Newmann, E., Idan, Y. & Pinkas, G. (1999). Discovery of fraud rules for telecommunications challenges and solutions, In *Proceedings of the 5th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Pages. 409--413. ACM Press.
- [18] Michael, H., Lambert, D., Pinheiro, C. & Sun, D. (2000). Detecting Fraud in the Real World, *Technical Report*, Lucent Technologies.
- [19] Murphy, K. (1998). A Brief Introduction to Graphical Models & Bayesian Networks, <http://www.cs.ubc.ca/~murphyk/Bayes/bnintro.html>
- [20] Russel, S. & Norvig, P. (2003). *Artificial Intelligence. A Modern Approach, 2nd edition, Prentice Hall Series Inc New Jersey 07458*.
- [21] Knorr, E. & Raymond T. (1999). Finding Intentional Knowledge of Distance-Based Outliers. *Proceedings of the 25th International Conference on Very Large Data Bases*, Pages 211 - 222, Morgan Kaufmann Publishers Inc.
- [22] Cozman, F. (2001). JavaBayes. Bayesian Networks in Java, University of Sao Paulo.

Isaac Olusegun Osunmakinde is a Ph.D. student in the Agents Research Group (ARG) of the Computer Science Department, University of Cape Town, South Africa. He has a M.sc. in Computer Sc. (UCT), a PgDS in Mathematical Sc. (Stellenbosch), a B.Sc. Hons in Computer Sc., & a HND in Computer Technology. He is a member of the IEEE Computational Intelligence Society. His current research interests in AI are probabilistic and mathematical modelling.

Anet Potgieter is a senior lecturer and head of the Agents Research Group of the Computer Science department, University of Cape Town, South Africa. She holds a Ph. D. (Computer Science) from the University of Pretoria (South Africa). Her current research interests include complex adaptive systems, distributed artificial intelligence, sensor networks and software engineering.