

Credit-based Network Management

Jilong Wang

Tsinghua National Lab. for IS&Tech
Network Research Center
Tsinghua University
Email: wjl@cernet.edu.cn

Dah Ming Chiu

Information Engineering Department
The Chinese University of Hong Kong
Email: dmchiu@ie.cuhk.edu.hk

John C.S. Lui

Computer Science & Eng. Department
The Chinese University of Hong Kong
Email: cslui@cse.cuhk.edu.hk

Abstract—Increasingly, a computer network administrator’s job is pre-occupied with user behavioral problems rather than physical failures of network and system components. A small number of malicious users can cause problems that affect a large number of users; more often, by not following proper procedures a user may let his/her system be used by malicious users; and there are various other misuses that all leave the network in a state of the *tragedy of the commons*. In this paper, we introduce the concept of *credit-based networking* - borrowing ideas from financial management and adapting them to network management. We first focus on a campus network by studying concrete scenarios of how credit-based network management can be applied. We then discuss how the concept is generally applicable to managing network behaviors as well by applying it to managing ISP peering relationships. We argue that the cascading effect of credit-based network management can enhance network management efficiency and improve the global network environment we all *live in*.

I. Introduction

Nowadays, computer network administrators are facing increasing challenges. The job of managing the network is no longer limited to learning about new technologies, upgrading software and hardware components, replacing broken equipments and such routine network management tasks. The new headaches are often caused by user behavioral problems.

For example, all sizeable networks get various security attacks on a routine basis. The methods for defense are mostly remedial: first find out where the security hole is, then download some patches to close the security hole. Currently, there is no strong deterrent against people who instigated the attacks because they are hard to catch; and no deterrent against people who open up security holes to be exploited either because they can only be accused of being negligent. Another type of behavioral problem is concerned with excessive use of network resources, to the extent causing service outage for other users. This type of scenario is quite common with the advent of various P2P content distribution applications. If the network capacity cannot be justifiably increased, then this becomes a network administrator’s nightmare.

Some network protocol designers are sensitive to these seemingly social issues. They design protocols to promote fair resource allocation to users, to the extent possible; they design systems to be difficult for malicious users to take advantage of; and in some protocol designs, mechanism design theory is used to build incentive for the software to do the socially acceptable thing. Yet in network management literature, there

is little discussion about how to effectively manage user behavioral problems. When faced with real users who do react to incentives, there is hardly any positive reinforcement for good behavior. In fact, authors in[8] suggest that how to manage user behavior is one of the most important challenges in network management.

The thesis of this paper is *credit-based networking*. In a nutshell, credit-based networking is to build incentives into ways people use the network so as to provide deterrence to bad behavior. This is similar to the use of credit in financial transactions to discourage and avoid bad outcome. We argue it can also be used to discourage and avoid bad network usage, and provide a more scalable solution to network management.

There are many technical challenges and issues in credit-based network management. For example, one basic question is whether there exists natural classification of users so that the *high risk* subclasses of users can be easily identified for more focused monitoring. For certain category of behaviors, a small number of individual users can cause significant negative impact to others. How to dynamically identify these users and keep their credit ratings is also a challenging problem. We discuss these issues in detail based on case studies.

In this paper, our discussions are based on the results and experience we have in managing a major university network in China (e.g., with approximately 35,000 users). We first show that based on past traffic analysis and trends, there is a need for credit-based network management techniques. We then discuss how credit rating might be kept and what might be in the administrator’s discretion to do to users with different credit ratings. To illustrate the utility of credit-based management, we apply the methodology to two specific problems: (1) security attacks based on ARP-spoofing, and (2) bottleneck caused by P2P traffics. We then discuss how the concept of credit-based networking can be generally applied to other scenarios of networking, for example ISP peering. Finally we conclude and suggestion several directions for further research.

II. User Behavior on a Large Campus Network

Increasingly, network management problems are caused by abnormal user behaviors rather than other abnormal events. Figure 1 and 2 show the percentages of different abnormal events in the network, collected from a major university in China. The percentage of those problems caused by user behavior is 55% in 2006, and increases to 63% in 2007.

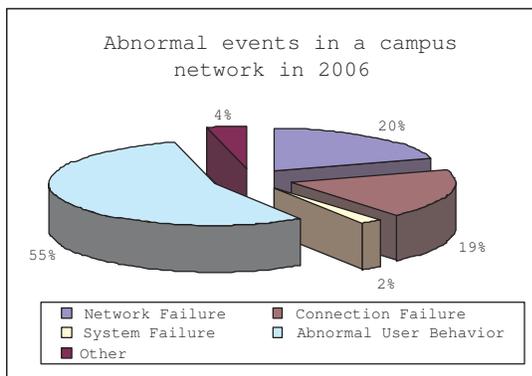


Fig. 1. Percentage of abnormal events in 2006

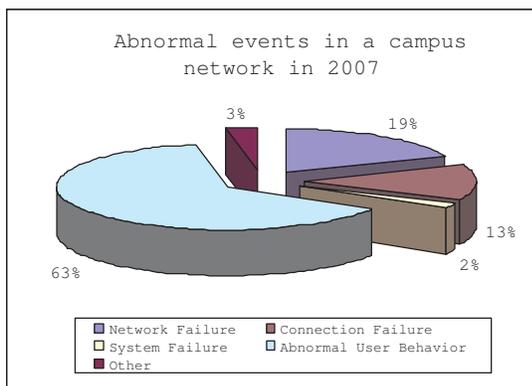


Fig. 2. Percentage of abnormal events in 2007

Network failure means the network abnormal events caused by the failures of routers, switches and other network equipments, while *connection failure* means those abnormal events caused by cable failure and other physical link problems. *System failure* indicates the events caused by the failures of DHCP, DNS and AAA type of services. Other abnormal events may be caused by *abnormal user behaviors*, such as ARP spoofing, DHCP spoofing, MAC Flooding, hacking, spamming, IP spoofing, bandwidth abuse, . . . etc. Many abnormal user behaviors are not intentional, e.g., their computers are infected with virus because of the vulnerability in the operating system or application software; or the computers are infected with Trojan horse because the users access malicious websites, emails or softwares. Although unintentional, these abnormal user behaviors need to be effectively managed since they can have a negative impact on network services to normal users.

Managing user behaviors is challenging for the following reasons. Firstly, the user population is usually quite large. For example, for the university campus network we study, the user population is 35,000, including staff, students and temporary visitors. Secondly, new problems continue to emerge with the rapid growth of the Internet. There is no long-term stable and effective tools. Each new problem often needs to be investigated manually and in a case-by-case basis, taking up considerable time. Thirdly, network management resources, operators and equipments, are always limited. Realistically,

we can monitor the whole behavior pattern of some users, or monitor certain behavior patterns of all users; but it is almost impossible to monitor all the behavior patterns of all users at all time. For these reasons, it is unreasonable to equally divide the limited network management resources to each user.

Let us take a deeper look at the abnormal events caused by user behavioral problems and see if they can be attributed to certain categories of users. Specifically, we extract all abnormal events known as *ARP spoofing*, generated during the period from April 2006 to December 2007. Each event record contains information of the user who caused the abnormal event. Three attributes are chosen for user information: user type, year of study, and department. The different user type are described in Table I. The year of study has the values of 1 to 6. “1” (“2”, “3”, “4”, “5”) means the year of study is 1 (2, 3, 4, 5), and “6” means the year of study is 6 or above. The values of the department range from 0 to 49, indicating the 50 departments that have users generating the abnormal events.

Based on the values of these three attributes, all users are divided into 3000 categories. There are a total of 1009 abnormal events in this dataset (corresponding to ARP Spoofing). The numbers of abnormal events for each user category are shown in Figure 3, which plots the number of user categories that have a certain number of abnormal events. Figure 3 shows that 2657 user categories (nearly 90% of all user categories) have no abnormal events (value “0” of the number of abnormal events). This result indicates that most of these

value	physical meaning
0	male undergraduate students
1	female undergraduate students
2	male master students
3	female master students
4	male Ph.D students
5	female Ph.D students
6	male staff
7	female staff
8	others (male)
9	others (female)

TABLE I
USER TYPE DEFINITIONS

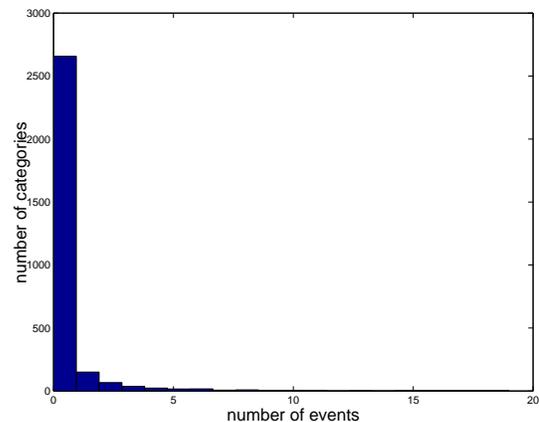


Fig. 3. the number of abnormal events

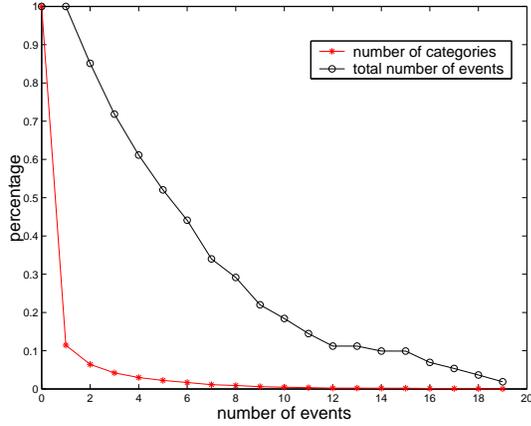


Fig. 4. CCDF curves for the number of categories and abnormal events

network abnormal events are generated by users from a small fraction of user categories. This ascertain our claim that we should not monitor all events from all users at all time, but rather, target it on small number of users.

To estimate the relationship of management costs and effect, we give out the complementary cumulative distribution function (CCDF) curves for the number of user categories and abnormal events in Figure 4. For the 3000 categories, each category produces different numbers of abnormal events, as shown in the x-axis of Figure 4. The curve with the asterisk marker is the CCDF curve of the number of categories with different number of abnormal events. Assume the number of abnormal event is e , and there are c categories that generate e events. Then the total number of events from these c categories are $N(e) = e * c$. The CCDF of $N(e)$ is shown by the curve with the circle marker. From Figure 4, one can find that given $e = 2$, there are about 6% of the total categories and 85% of the total events. In other words, if we monitor those categories which produce more than two abnormal events, then we only need 6% of the management resources to monitor 85% of the abnormal events. We call the users in these categories *questionable users*. If we only monitor the behavior of the questionable users, we can greatly reduce the management costs while keep good management effect.

The next question is whether we can find *which* user categories are generating the majority of these abnormal events, at least for the network management of a large campus network. We define *class* as the categories with the same value in a certain attribute. For example, the “user type = 0” class includes all 300 categories that have the same value “0” for the attribute “user type”. In addition to the three attributes, we define class “gender = male” as the 1500 categories with “user type = 0,2,4,6,8”, and so does “gender = female”. These classes are shown in the first and second column in Table II. All categories have the same value (the “classes” column) for an attribute (the “attributes” column). The third column “ n ” means there are n categories in the class. For the attribute “department”, we only select the top 10 departments which generate the largest number of abnormal events and omit the

attributes	classes	n	$\bar{X} \pm S$	Z, P
gender	male	1500	0.47 ± 1.65	4.76
	female	1500	0.21 ± 1.00	,0.000
user type	0	300	1.43 ± 2.93	402.51 ,0.000
	1	300	0.77 ± 2.02	
	2	300	0.15 ± 0.89	
	3	300	0.09 ± 0.50	
	4	300	0.69 ± 1.66	
	5	300	0.13 ± 0.51	
	6	300	0.06 ± 0.26	
	7	300	0.02 ± 0.15	
	8	300	0.00 ± 0.06	
	9	300	0.01 ± 0.11	
department	2	60	1.32 ± 2.95	225.67 ,0.000
	12	60	1.05 ± 2.05	
	20	60	0.98 ± 2.45	
	3	60	0.92 ± 2.88	
	10	60	0.92 ± 3.14	
	41	60	0.83 ± 2.12	
	28	60	0.77 ± 2.29	
	45	60	0.75 ± 2.72	
	36	60	0.73 ± 2.64	
	27	60	0.65 ± 1.64	
year of study	1	500	0.38 ± 1.18	108.68 ,0.000
	2	500	0.58 ± 1.85	
	3	500	0.64 ± 2.12	
	4	500	0.32 ± 1.21	
	5	500	0.05 ± 0.37	
	6	500	0.04 ± 0.23	

TABLE II
AVERAGE NUMBER OF EVENTS FOR DIFFERENT CLASSES

other 40 departments.

There are number of ways to discover “questionable” users. One simple method is to look for the classes with large number of abnormal events. In one class, there are n categories each generates e_i abnormal events. We define \bar{X} as the mean of the e_i for the n categories, and S as the sample standard deviation. $\bar{X} \pm S$ of each class is shown in the fourth column of Table II. \bar{X} of each class is an intuitive indicator of whether it contains a significant number of questionable users. For example, \bar{X} of the class “gender = male” (0.47) is greater than that of the class “gender = female” (0.21), so the class “gender = male” contains more questionable users.

We can apply non-parametric tests[5] to further determine if there is a difference between different classes. For the attribute with two classes such as “gender”, we use Mann-Whitney U test. The null hypothesis is that the two samples are drawn from a single population, and therefore their probability distributions are equal. For the attribute “gender”, the null hypothesis means there is no difference between the classes “gender = male” and “gender = female”. We use SPSS, a statistical analysis software, for data analysis, which outputs Z values or P values to determine significance. When using Z values, if Z is less than or equal to the critical Z value of 1.96 ($P \leq 0.05$), then we can assume that the null hypothesis is correct and there is no difference between the two classes. On the other hand, if Z exceeds 1.96, then we reject the null hypothesis. It is even more convenient to test using the P values. If the P value is low, then there is difference between the two classes. For the attributes with more than

two classes such as “user type”, “department” and “year of study”, Kruskal-Wallis H-Test is used, which is an extension of the Mann-Whitney U test to multiple samples. Z values and P values are shown in the last column of Table II. From them one can find that there is difference between different classes for these attributes.

Base on the statistical results of abnormal user behaviors, we conclude that given the limited network management resources, we can increase the priority of resource usage for questionable users, and establish a more effective and scalable management policy. Here are some examples:

- Give training to freshman on network security, how malicious users exploit other users’ machines for their attacks, and network usage guidelines to avoid security holes.
- Provide special support for certain classes of users. For example, provide technical consultation for retired teachers, helping certain users to configure systems correctly, and clean virus and Trojan horses.
- Pay more attention to monitoring the LANs or departments that contain questionable users.

III. Credit-based Network Management

Through the analysis and discussion in the last section, we have argued for the following management philosophy - there are a large number of users and only limited management man power; but the questionable users represent only a small percentage of user population; so it is important to separate users into different classes and focus the management resources on the questionable users. This is the basic idea of credit-based network management.

A more fully developed practice of credit-based network management would include the following aspects:

- 1) Classify users according to an analysis of historical behavioral patterns.
- 2) Apply focused monitoring according to user classifications (paying more attention to questionable users); assign credit scores to users according to observed behavior; apply different management policies (service) to users with low credit.
- 3) Dynamically adjust the credit score for individuals, and user classes as well. Although a well-designed credit system can assign credit scores for most users appropriately according to their past behavior, in a small number of cases, human intervention through re-adjusting scores for some individuals may be necessary.
- 4) Rely on the psychological effect of credit scores to deter bad user behavior. Once the users care about their credit scores, behavioral problems will accordingly decrease.

The process of credit-based network management can therefore be organized into three levels of activities: (a) network/user behavior measurement; (b) network/user credit evaluation; and (c) network/user behavior control, as shown in Fig. 5.

The first and more fundamental level is network/user behavior measurement. Network measurement is an active research

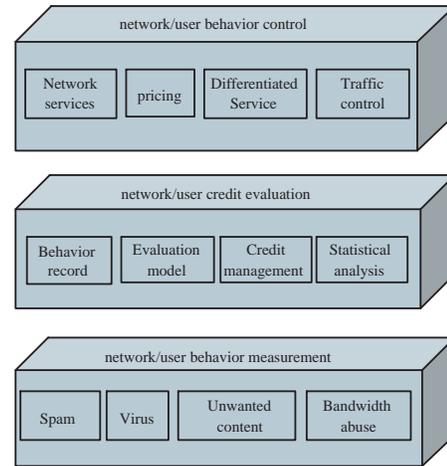


Fig. 5. Infrastructure Design of Credit-based Network Management

area. There are lots of research results that can be applied, especially in the traffic analysis and identification area. The second level is the evaluation of network and user behaviors based on the measurement results, which is essentially a kind of classification problem in the study of pattern identification. One can use methods such as discriminate analysis, classification tree and SVM, etc. At the top level, we control network/user behaviors based on the evaluation results. This may be done through providing different levels of services content, or applying different pricing schemes or different traffic control. Note that even if we do not apply any behavior control policies, the credit system itself can still have an effects on users, because it is human nature for users to want to get high credits. Since credit-based network management is about applying traditional management ideas to network management, it comes with a solid theoretical and technological foundation.

There is, however, some notable differences between applying credit scoring to financial lending versus network management. In traditional credit systems, lenders (such as banks and credit card companies) use credit scores to evaluate the potential risk of lending money to different consumers. A credit score is derived from a person’s past actions, based on a probability model that can be used to predict the likelihood that the person will pay back debts in a timely manner. When applying traditional credit scoring to network management, credit-based network management first analyze network abnormal behaviors, and then classify users. The classification results are then used to determine network management policies. There are several differences between credit-based network management and the traditional credit scoring system.

- **Classification results:** the first difference is the classification results that are required. In traditional credit scoring system that is used by banks, the most important thing for classification results is its accuracy. However, in network management, we also require that the classification results to have physical correspondence in the network. For

example, if the classification results indicate that the questionable users are from several departments or several dormitories, the results are very useful in network management, since we can offer more management resources to monitor these departments or dormitories. However, banks can use credit scores to evaluate each customer and determine who qualifies for a loan.

- **Credit index:** in traditional credit scoring systems used by banks, credit indexes are used to calculate personal credit score. To get accurate results, a bank may use many credit indexes each containing very personal information about the customer. The credit index may include the customer's occupation, marriage status, income, asset, etc. In network management, we need to consider their physical meaning when choosing credit indexes (we call them attributes in network management). It is better that the users with the same attributes are situated in the same local network or place. Based on this consideration, IP address, department and residence address are appropriate attributes for network management.
- **Scoring model:** in traditional credit score system used by banks, there are two kinds of errors in classification that needed to be considered. One is false negative. The banks mistakenly classify the customers who would not pay debts in time and give them high credit. The other is false positive, which means the banks mistakenly classify those customers who would pay debts in time and give them low credit. False negatives usually bring more losses to banks due to bad debts. This should be taken into consideration when choosing a credit score model. However, in credit-based network management, the model selection should first consider the physical meanings of the classification results. Therefore, some credit score models that are widely studied and used in traditional credit scoring systems, such as neural network models[7], [10], may not appropriate for network management.

IV. Case Study

In this section, we present a case study of how to apply credit-based network management in the campus network we introduced earlier in Section II. Due to the lack of space, we focus on two issues: (a) user classification¹; (b) assigning and adjusting individual user scores.

A. User Classification

We still use the same data as in Section II with 1009 abnormal events.

1) **Attributes selection:** Every record has four attributes: user type, year of study, department, and address. The first three attributes have been discussed in Section II. We divide the addresses of all users into 51 areas. To avoid the difficulty brought by the large number of categories, we use two stage analysis. In the first stage user type, year of study and

department are used, while in the second stage user type, year of study and address are involved.

2) **Model Selection:** A wide range of statistical classification methods has been applied in credit scoring[9], such as discriminate analysis, linear regression, logistic regression, logistic regression, classification trees and neural networks. Based on the requirement of credit-based network management that the classification results should have physical meanings, we choose classification tree as the classification model, since its results have intuitive meaning.

Classification tree ([2], [3]), also known as decision tree, is widely used in many areas. Applications of such method in credit scoring is described by [6]. An important feature of this method is its capability to break down a complex decision-making process into a collection of simpler decisions, thus providing a solution which is often easier to interpret. A classification tree is a flowchart-like tree structure, where the internal node denotes a test on an attribute, the branch represents an outcome of the test, and the leaf node has a class label.

Classification tree method has several advantages. The construction of classification tree does not require any domain knowledge or parameter setting. It can handle high dimensional data. The representation of classification result in tree structure is intuitive and easy to be read by humans. In general, classification tree has good accuracy, and its learning and classification steps are simple and fast.

Many algorithms such as ID3, C4.5, CART are developed for learning classification trees. These algorithms adopt a greedy approach in which classification trees are constructed in a top-down recursive manner. The algorithm starts with a set of training tuples with their associated class labels, a list of attributes and an attribute selection method. The tree starts as a single node (the root node), representing all the training tuples. The algorithm uses the attribute selection method to determine the **splitting criterion**. The splitting criterion tells us the best attribute at the root node to partition the tuples into individual classes. It also tells us which branches to grow from the root node with respect to the outcomes of the chosen test. The entire process is repeated using the training tuples associated with each descendant node to select the best attribute to test at that point in the tree.

The attribute selection measure is for selecting the splitting criterion that "best" separates a given data partition, D , of class-labeled training tuples into individual classes. Popular attribute selection measures include, information gain, gain ratio, and gini index.

3) **Result analysis:** We use the classification tree algorithm[4], [1] to analyze the above data. The classification results are shown in Figure 6 and 7. In Figure 6, the list of attributes includes user type, year of study and department. In Figure 7, the list of attributes includes user type, year of study and address. In each node of the tree, *Mean* represents the average number of abnormal events from all the categories in the class, *Std. Dev.* represents the sample standard deviation, n is the total number of categories in the class, and % indicates

¹In Section II, we described what user classification means and methods. Here, we try to performance the classification.

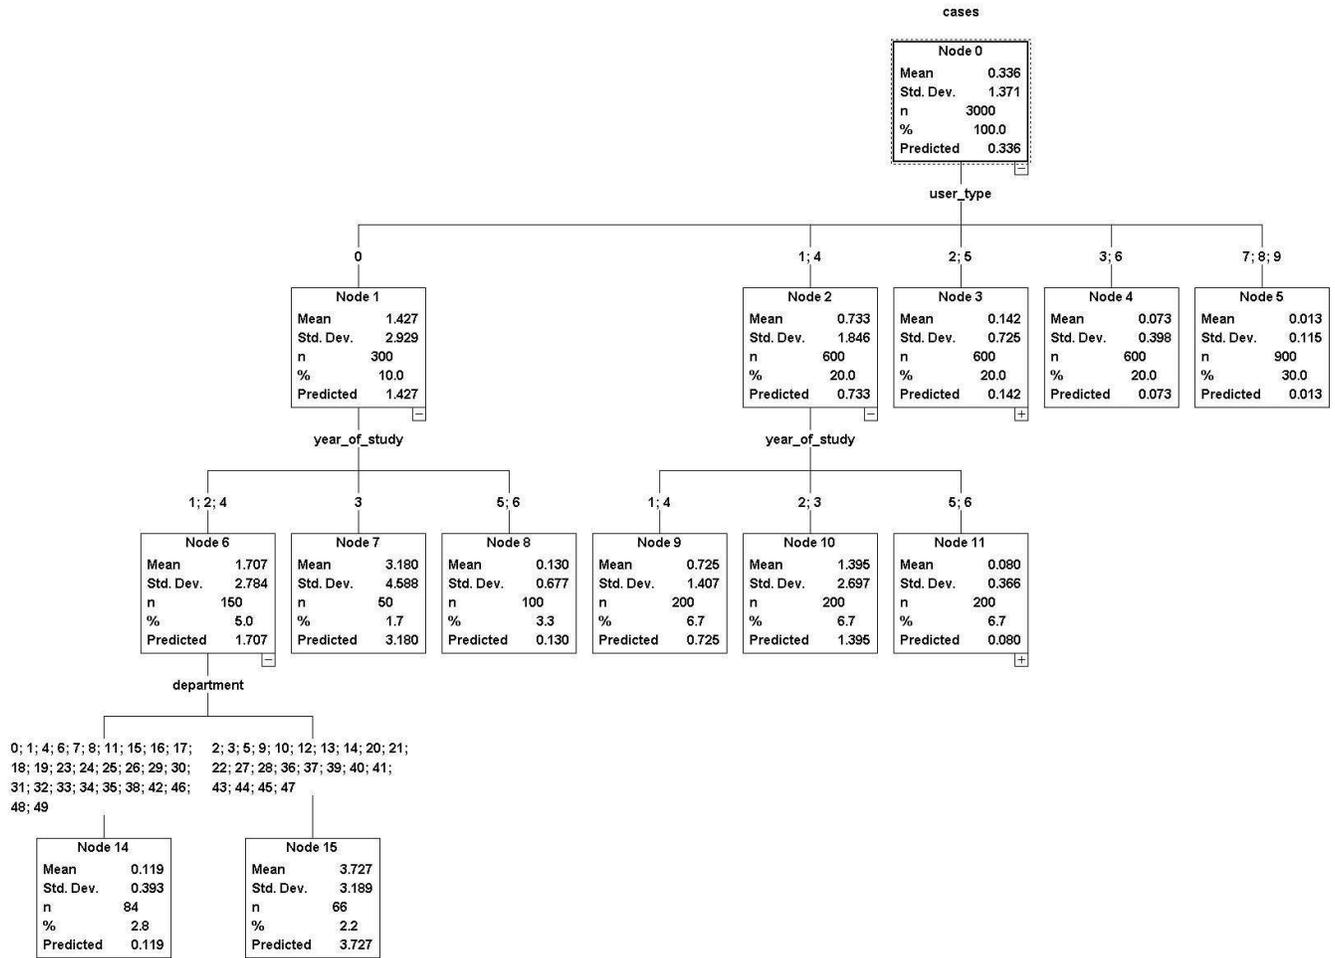


Fig. 6. Classification tree: user type, year of study and department

the percentage of n in the total number of categories.

From Figure 6, we find that male undergraduate students (node 1) have the highest probability to generate abnormal events for the attribute “user type”. Among all male students, those third-year students (node 7) produce more abnormal events. For the male students whose year of study equals to 1, 2, 4, some departments (node 15) generate more abnormal events than others. In addition, node 10 also produce large number of abnormal events, which presents female undergraduate students, male phd students with year of study of 2 and 3. From Figure 6, we also find that user type and year of study are important attributes, while the attribute department only affects some users. From Figure 7, one can find that address is a more important attribute than year of study. For male undergraduate students, network abnormal events centralized in 11 areas (node 6). It’s mean value is 6.4, which is much higher than other areas. Another centralized area is represented by node 8.

Based on the above results, we can identify those question-

able users and accordingly design and manage the campus network, which can better utilize the management resources. For example, we set up a new network equipment that provides good protection to ARP spoofing in a centralized area of the identified questionable users.

B. Individual Credit Scoring

The reason for classifying users is that some classes of users have much higher probability than other classes for abnormal behavior. However, there may be some individual users who generate more abnormal events. We discuss how to assign and adjust individual credit scores for these users in this subsection. From the management records in the campus network, we find that 7.1% of the users repeatedly generate abnormal events, which account for 14.3% of all the abnormal events. However, user classification is not sufficient for this case, since these users do not have significant statistical properties. Therefore, in addition to classifying users, we also need to monitor and manage individual users who have a

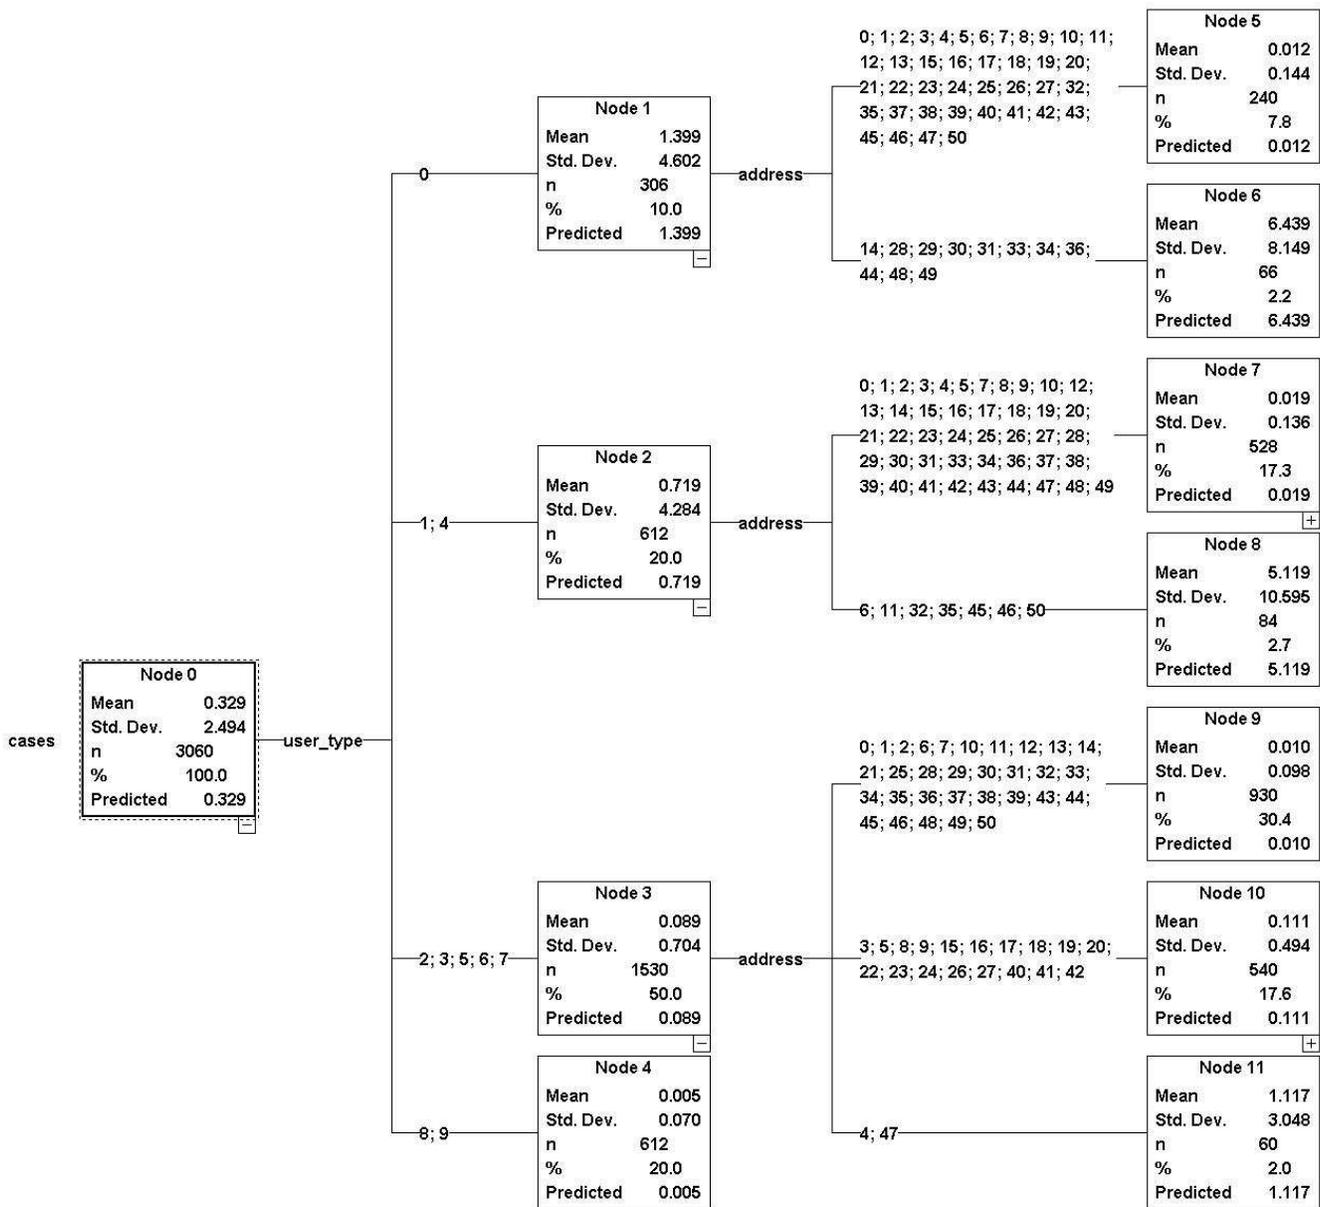


Fig. 7. Classification tree: user type, year of study and address

history of producing abnormal events.

If we monitor and manage every user who have records of abnormal events, the problem is that the number of monitored users will keep growing. Our solution is to use individual credit score to identify questionable users, and to memorize past “bad” behavior of users. Individual credit scores are adjusted according to users’ behaviors. When a user generates an abnormal event, his credit score will decrease. After the user’s credit score falls below a threshold, the user becomes

a questionable user and will receive special monitoring and management attention. On the other hand, if a user does not generate any abnormal event, his credit score will increase over time. Through this process, a questionable user can gradually return to become a normal user.

We calculate individual credit scores based on the following considerations:

- Different kinds of abnormal events have different effects to credit scores, and more serious events cause higher

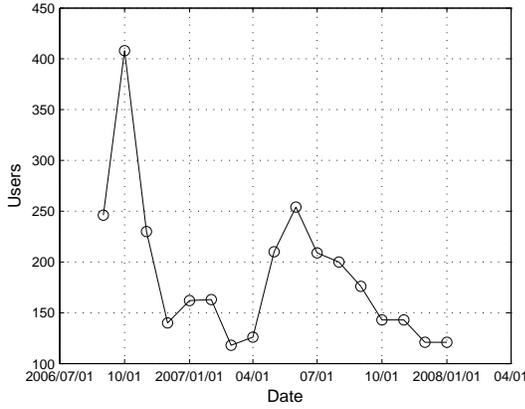


Fig. 8. the number of questionable users

deduction from credit scores.

- Deliberate bad behavior should get more deduction than non-deliberate ones.
- For users who generate abnormal events with high frequency and or within a short time interval, the credit score should see more deduction.
- The amount of credit deduction also depends on the current value of the credit score, since a low value implies that the user is a repeat offender.

Based on the management data on the campus network, we have designed the following procedures for calculating individual credit scores, on an experimental basis:

- All users have an initial value of 60 for individual credit score. (Another possibility is to assign the initial individual credit score based on user classification: give lower credit value to those users in the identified classes that have larger number of abnormal events. The disadvantage of this method is that good uses in these classes are penalized).
- When a user is detected to have bad behaviors, the credit score is deducted as follows:
 - 1) $newcredit = curcredit - 10 * w$ for the first detected event,
 - 2) $newcredit = curcredit - (curcredit * 0.5 + 10 * w)$ for the second and later detected event,
 where $curcredit$ is the current credit score, $newcredit$ is the new credit score, and w is the weight for different abnormal events. For example, w for ARP spoofing is 1.2 because of its severity, w for DHCP spoofing is 1.5 since this kind of events is more likely to be intended behaviors, and w for most other abnormal events is 1.1.
- A user becomes a questionable user when his credit score is below 50.
- The credit score is increased by 1 each month if the user is not associated with any abnormal events in that month.

Based on the network management records of 2006 and 2007, we can do a calibration exercise. The number of questionable users in each month is shown in Fig. 8. According to the above description of how credit score is calculated,

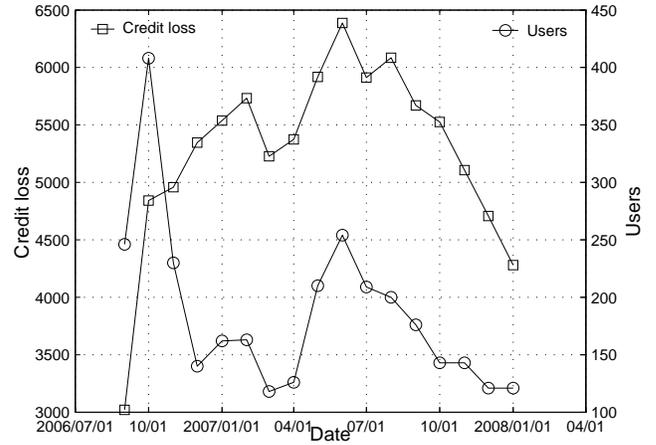


Fig. 9. the amount of cumulative credit loss

questionable users have the credit value less than 50. Since every user's credit score increases by 1 every month, it takes a user who generates one abnormal event about 2 months of normal behavior to be removed from the list of questionable users. From this figure, one can find that the highest number of questionable users is about 400. The reason for significant reduction in the number of questionable users in the second half of 2007 is due to an educational campaign plus various new network management measures. Still, the lowest number of questionable users stabilizes around 120, which is about 0.03% of the total population.

To represent the long term situation of the campus network, we introduce another variable, the cumulative credit loss. Assume there are totally N users in the network, and m users generate abnormal events in the j^{th} month, then the total amount of credit deduction of all the m users is $credit_loss = \sum_{i=1}^m (curcredit_i - newcredit_i)$. For the remaining $N - m$ users, assume there are k users whose $curcredit < 60$, then the total amount of credit increase of the k users is $credit_incr = k * 1$. So the credit loss in the j^{th} month is $dlt_credit_j = credit_loss - credit_incr$, and the cumulative credit loss of the j^{th} month is $cum_creditloss_j = cum_creditloss_{j-1} + dlt_credit_j$. Figure 9 shows the amount of cumulative credit loss every month.

We introduced individual credit scoring to the campus network management recently, and have already achieved some effect. Firstly, as mentioned earlier, we applied some targeted actions to the questionable users: special training sessions, inspection of certain computer systems, and for severe cases, warnings were sent. These actions reduced the number of questionable users as well as the amount of credit loss, as shown in Figure 8 and 9. In addition to guiding network management, individual credit scoring also influence user behavior automatically, as users are all wary of bad records.

In the long run, we are also considering to apply credit scoring to the following situations:

- The users with low credit will be constrained when they

apply for some advanced services.

- Provide low QoS to the users with low credit.
- Charge more to the users with low credit.

V. Traffic Control of Campus Network Internet Access Link

The Internet access link of a campus network can easily become the traffic bottleneck. The access link will be saturated on peak time, and causes packet loss and long delay. The solutions of this resource constraint include:

- Upgrade the network equipments and increase the bandwidth. This is an effective method but requires high economic cost.
- Provide differentiated services based on application type. Based on the analysis on the outgoing traffic of the campus network, we find that the amount of traffic that is detected or suspected to be P2P applications counts for a big proportion. But we encounter problems when we try to restrict the bandwidth of P2P traffic. If we restrict the bandwidth of the detected P2P traffic to 20M, the congestion situation does not improve much, while suspected P2P traffic increases. However, we can not unreasonably restrict the suspected P2P traffic, since it will affect normal network usage. Generally, the solution of providing differentiated services to different applications has several problems:
 - The cost and accuracy of application identification. Accurate application identification usually needs more than examining IP and TCP/UDP headers of packets. And it becomes more and more difficult because of the growing complexity of new network application protocols.
 - It is hard to define appropriate rules for differentiated services. It is the abuse behaviors that need to be constrained instead of some certain applications. For example, by analyzing the traffic, we find that a large proportion of all traffic is P2P traffic and suspected P2P traffic. But to achieve resource control by simply limiting P2P traffic might not be supported by users.
 - Constraining or prohibiting new applications will be an obstruction to novel technologies.
- Allocate traffic bandwidth by users: A “fair” solution is to allocate bandwidth by users, setting rate limits by charge standards. Experiments have been done on TUNET, setting up queues by IPs for VIP users and normal users respectively. It turns out to solve the bottleneck issue efficiently. But this solution also has some problems. For example, the static allocation can not meet the requirements of some outburst communications.
- Use *credit-based network management*: The root cause of the problem is the abuse behaviors, so we apply the credit-based network management, and try to solve the problem from the root. The main steps are: a) Calculate and record the user credits according to their traffic statistics. If a user is recorded as occupying high bandwidth

for a long time, then his credit will be deducted by some amount. A user will be classified as abuse user when his credit becomes lower than a certain threshold. b) Set up an unrestricted queue for normal users, and a restricted one for abuse users. So we can restrict abuse users and at the same time keep the normal users unaffected. c) There is still a problem. Normal users can have abuse behaviors and cause bandwidth problems before it is classified as abuse user. To solve this problem we can further use classification of user credit, limiting bandwidth for user groups that have high possibility to have abuse behaviors.

VI. Generalization

In addition to campus network, the credit-based network management method can be generalized and applied to various network management tasks. In order to achieve this, five entities, i.e. Who, Whom, What, How, and Expectation, need to be defined for each scenario. We illustrate three different scenarios here.

A. Scenario 1: Credit-based Network Management on High-level Backbone

Backbone network does not generate traffic itself. All traffic comes from the customer network. Traffic classification and filtering are usual solutions to constrain unwanted traffic. These solutions are able to reduce unwanted traffic on backbone network while they are incapable of constraining unwanted traffic from the source, i.e., the customer network. In addition, traffic filtering might increase the burden of the network, and has the risk of impacting legal communications. If we introduce credit-based network management to backbone network, we can constrain unwanted traffic from the source customer network.

Who: Internet backbone provider

Whom: Customer network

What: the number of unwanted traffic, bandwidth usage, etc.

How: by service content, or Quality-of-Service (QoS) policy

Expectation: Reduction of unwanted traffic

The result of credit scoring can be used for adjusting price or QoS to the customer network, which encourages customers to enhance their self-management. Credit scoring might not change the traffic situation immediately, but it could have a long term impact on customers generating unwanted traffic.

B. Scenario 2: Credit-based Network Management between peers

Can credit-based network management work for peers? Figure 10 illustrate relationship between peers. The objective of establishing peer connection is to optimize routing and improve communication performance. But the unwanted traffic flooding between peers could also cause serious problems. So it is meaningful to let peers supervise and coordinate with each other.

Who: ISP

Whom: Peer

What: Unwanted traffic

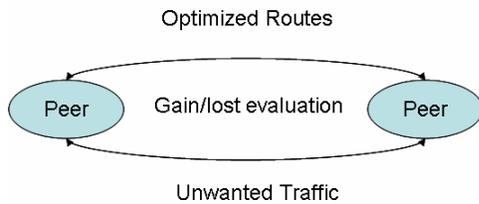


Fig. 10. Peer relationship

How: Traffic engineering

Expectation: Improve peer's value

C. Scenario 3: Credit-based Network Management on Intrusion Detection

Intrusion detection usually requires deep traffic analysis such as examining packet payload, which has serious resources consumption problem. We can use credit-based network management to improve the efficiency of intrusion detection systems. Scanning is usually a an indication of network attacks, which is easy to detect. We can classify IP blocks by detected scanning behaviors, and only implement deep traffic analysis on the identified IP blocks. Thus we can have more resources and provide more complicate analysis.

Who: Intrusion detection system

Whom: IP blocks

What: Unwanted scanning behavior

How: Deep traffic analysis

Expectation: Increase analysis efficiency

VII. Conclusion

In this paper, we propose the idea of *credit-based network management*. In the broadest sense, the idea is to advocate *social responsibility*. Since a network is a place shared by many, a small number of users can easily cause inconvenience to a large number of users. Credit-based networking lets network administrators set up policies to encourage good behavior and deter bad behavior, thus foster a good networking environment for all users to share.

We also study and discuss various technical challenges in implementing this idea in a real network setting: the campus network of a large university. This includes how to classify users based on past behavior patterns; how to design and adjust credit scoring for individuals; as well as how to use the idea to solve a real-life problem - managing congestion in the network exit link.

There are many directions for further study. Although based on real-life network management considerations, many of the ideas are yet to be implemented and tested. The explorations of user classification and credit adjustment mechanisms are preliminary, and can benefit a more rigorous and formal approach. Fortunately, the framework is quite robust, in the sense that different policies, algorithms can be experimented without irrevocable consequences. From a more global perspective, the effect of practicing credit-based networking can cascade, and incrementally make the global network a better place. Yet there

is no need for close coordination for different networks. When credit-based networking become more widely practiced, it will then be time to consider a consistent way of maintaining credit so that a user can carry his credit from one network to another as well.

Acknowledgement: this research is supported in part by the NSFC-RGC Grant: N_CUHK414/06.

REFERENCES

- [1] D. Biggs, B. de Ville, and E. Suen. A method of choosing multiway partitions for classification and decision trees. *Journal of Applied Statistics*, 18, 49-62, 1991.
- [2] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA, 1984.
- [3] J. Han and M. Kamber. *Data Mining: Concepts and Techniques*. Morgan Kaufmann Publishers, 2001.
- [4] G. Kass. An exploratory technique for investigating large quantities of categorical data. *Applied Statistics*, 29:2, 119-127, 1980.
- [5] E. Lehmann. *Nonparametrics: Statistical Methods Based on Ranks*. San Francisco: McGraw-Hill, 1985.
- [6] P. Makowski. Credit scoring branches out. *Credit World*, 74, 1985.
- [7] M. Odom and R. Sharda. A neural network model for bankruptcy prediction. In *Proceedings of the IEEE International Joint Conference on Neural Networks*, 1990.
- [8] P. J. Schonwalder, M. Burgess, O. Festor, G. Perez, R. Stadler, and B. Stiller. Key research challenges in network management. In *IEEE Communications Magazine*, volume 104-110, 2007.
- [9] L. C. Thomas. A survey of credit and behavioral scoring: Forecasting financial risk of lending to consumers. *International Journal of Forecasting*, (16): 149-172, 2000.
- [10] D. West. Neural network credit scoring models. *Computers and Operations Research*, 27: 1131-1152, 2000.