# Analysis of Security Policy in Practical Internet Coordinates

Xiaohan Zhao, Xiaoxiao Song, Xiao Wang, Yang Chen, Beixing Deng, Xing Li
Tsinghua National Laboratory for Information Science and Technology
Department of Electronic Engineering, Tsinghua University, Beijing 100084, China
*homeisland03@gmail.com*

## *Abstract*

*Network Coordinate (NC) System is an effective mechanism to predict network delay with limited measure overhead. As one of the representative NC systems, Practical Internet Coordinates (PIC) has proposed a security policy based on triangle inequality to defend malicious nodes in the system. However, there is a natural phenomena that nodes may violate triangle inequality in Internet. Thus, the performance of PIC security policy is worthy to be well researched under different attacks. In this paper, we analyze the security policy in PIC in three real network delay data sets and compare PIC with security to without security under four typical attacks. The experimental results in this paper demonstrate that PIC is vulnerable to attacks while more TIVs will cause higher relative error of PIC. Moreover, under attacks by more than 40% malicious nodes, the performance of PIC with security policy could barely be better than PIC without security. Even Colluding Isolate attack will result in worse performance in PIC having security policy.*

## 1. Introduction

In recent years, network coordinate systems, which can accurately estimate network delay or round trip time (RTT) with low overhead, have been proposed to support the application overlays that benefit from topology-awareness. Up to now, there are several widely studied network coordinate systems: GNP [1], NPS[2], PIC [3] ,Vivaldi [4] and Pharos [5], where GNP is a centralized coordinate system because it requires several settled nodes as its landmarks while the rest rely on all the nodes in systems, thus they are regarded as decentralized.

However, since a network coordinate system must run a long time and it is consisted of a large number of nodes, it would attract hackers to attack the system. Moreover, decentralized network coordinate systems [2-5] ideally assume that nodes in them act honestly, in other words, nodes in such system will respond others with its true information, such as its coordinates, the RTT between the probing nodes and themselves and so on. But such assumption cannot be realized in real network and it would make network coordinate systems vulnerable to malicious attacks.

Although it has been verified by Kaafar et al. that Vivaldi and NPS are sensible to attacks from inside malicious nodes [6][11], for PIC, an representative network coordinate system which has been applied in an important P2P system---Pastry [7], the effect of attacks on it has not been studied yet, which leads to the study of attacks on network coordinate systems incomplete.

Furthermore, PIC has proposed a security policy relying on triangle inequality to detect malicious nodes. However, numerous research works have proved that there exist Triangle

Inequality Violation (TIV) [12] in Internet as a result of the Internet structure and routing policy. Therefore, this triangle inequality based security policy in PIC faces great challenge. Whether the policy can secure PIC efficiently or not is a significant problem worthy to be researched.

In this paper, we study the performance of security policy in PIC under different attacks in various network conditions. We utilize three data sets and apply four types of attack on PIC, which require fewer restrictions but are more typical than the attack proposed in [3]. The results show that PIC is vulnerable to attacks as well as more TIVs can cause higher error in NC computation. Moreover, when the percentage of malicious nodes is larger than 40%, PIC with security performs barely better than without security and even worse in Colluding Isolate Attack.

The rest of the paper is organized as follows. Section 2 presents an overview of PIC, TIV and classification of attacks. How the simulation is set up and the results of experiment are demonstrated in section 3. Finally, in section 4, we conclude this paper.

## 2. Related Work

### 2.1. PIC Overview

PIC is a decentralized coordinate system. Because there are no fixed nodes called infrastructure nodes to assist coordinate computation and every node can be chosen by other nodes as their landmark, which is used to compute the coordinates of other nodes.

In PIC, each node selects L nodes as its landmarks to compute its d-dimensional coordinates (L>d) in Euclidean space. When the number of nodes already in the system N is smaller than the number of landmarks L, the joining node selects all nodes as its landmarks, collects measured distances between all pairs of these nodes, which constructs a N×N matrix and uses a global optimization algorithm to compute a new set of coordinates for all nodes.

When nodes in the system are more than L, a different method is used to compute coordinates: the joining node chooses L landmarks from all the already-in nodes and after getting information from its L landmarks, it computes its corresponding coordinates by exploiting an optimization algorithm to minimize the error between the predicted distance and the measured distance. The target error function is the sum of the squares of the relative errors:

$$e = \sum_{i=1}^{|L|} \left(\frac{d_i^m - d_i^p}{d_i^m}\right)^2 \qquad (1)$$

Where $d_i^m$ is the measured distance between the joining node and its ith landmark and $d_i^p$ is the predicted distance between them.

To select landmarks, the author of [3] presents three strategies: 1) landmarks are chosen randomly; 2) only the closest nodes can be selected as landmarks for the joining node; 3) some of landmarks are picked randomly while the rest are picked from the closest ones. The results in [3] show the third strategy performs best and it is used in this paper.

Considering the existence of malicious nodes, PIC proposes a security policy based on triangle inequality. The main idea is that each node uses triangle inequality to verify every landmark and rejects the one which violates the triangle inequality mostly. Specifically, after a node *n* receives information from its landmarks, it computes the two metrics below for every landmark *i*:

$$upper_i = \sum_{j=1}^{|L|} \begin{cases} d_i^m - (d_j^m + d_{i,j}^p), & if\ (d_j^m + d_{i,j}^p) < d_i^m \\ 0 & otherwise \end{cases} \tag{2}$$

$$lower_i = \sum_{j=1}^{|L|} \begin{cases} (d_j^m - d_{i,j}^p) - d_i^m, & if\ (d_j^m - d_{i,j}^p) > d_i^m \\ 0 & otherwise \end{cases} \tag{3}$$

Where $d_i^m$ is the measured distance between node *n* and landmark *i*, so is $d_j^m$. And $d_{i,j}^p$ is the predicted distance between landmark *i* and landmark *j*. Then, the maximum values of both metrics are found out, the corresponding node is eliminated and the joining node computes its coordinates with the remaindering landmarks until the above process is repeated for required times or the average relative error between the joining node and the remaindering landmarks excesses the settled threshold.

As mentioned in [3], PIC has been used in Pastry [7], a generic, scalable and efficient substrate for peer-to-peer applications. The usage of PIC in Pastry can largely reduce the control traffic. Studies in [3] also show that PIC works well in a churn environment where nodes join and leave the overlay continuously. In other words, PIC has demonstrated its usefulness in Proximity-aware P2P overlays.

### 2.2. Triangle Inequality Violation (TIV)

In Internet, because of the natural feature of Internet structure and routing policy, the measured delays between nodes do not always match triangle inequality. For example, if there are three nodes A, B and C and even A, B and B, C are very close to each other, the delay between A and C may be not very small. That is, AC>AB+BC, which is a Triangle Inequality Violation, named TIV for short.

If an edge AC causes TIV, its triangulation ratio of the violation can be calculated with the definition T=*d*(A,C)/(*d*(A,B)+*d*(B,C)). Based on the definition of TIV, it is simply to judge whether a edge causes TIV: if T>=1, the edge is obviously against triangle inequality; otherwise, the edge adapts triangle inequality.

In order to evaluate the severity of TIV of an edge, [13] defined a metric, TIV severity, as the equation (4) considering the triangulation ratio of the violation and the number of TIVs caused by the edge. If TIV severity metric value is 0, it means the edge induces no TIV. Otherwise, the edge violates triangle inequality more when the value of this metric is large.

$$S = \frac{\sum d(A,C)/(d(A,B)+d(B,C))}{|N|} \tag{4}$$

The analysis in [13] demonstrates that although more severe TIV tends to result from long edge, there is no evidence to indicate the relationship between edge length and severe TIV. Thus, whether an edge will cause severe TIV cannot be indicated by the length of edge. In addition, since even close-by edges are different in TIV severity, it is hard to predict TIV based on the proximity of two nearby edges. Thus, TIV is a real but complex existence in Internet

However, PIC aims to embed nodes into a low dimension Euclidean space. Thus, the coordinates of every node must accord with triangle inequality. If delays between any two nodes obey triangle inequality ideally, the computed coordinates in PIC will perfectly match the real network and have low relative error. But the existence of TIV in real network makes it impossible that PIC predicts all network delays accurately and will cause high error between some edges. Moreover, since PIC security policy defines malicious nodes as nodes that violate triangle inequality mostly, according to our tuition, it is possible that security policy would falsely detect nodes in severe TIV edges as attackers and the security policy may not secure PIC well, which will be analyzed in section 3.

### 2.3. Attack Classification

[6] classifies attacks on network coordinate systems into four classes:

(1) Isolation: Malicious nodes select several nodes as their targets and then inveigle themselves into a remote area. These targets seem to be isolated from other nodes so that they would probably choose malicious nodes as their neighbors because the malicious ones are their closest nodes in the remote zone.  Thus, malicious nodes can play tricks on these targets.

(2) Repulsion: In order to reduce the consumption of its resources, e.g. bandwidth, a malicious node provides other nodes with false information either by forging coordinates or delaying the probes to pretend its position is rather far away.

(3) Disorder: The aim of this attack is to cause high error or even non-convergence in coordinate systems. In order to realize the attack, malicious nodes provide fake information to others.

(4) System control: In this attack, malicious nodes try to be in higher hierarchy to influence as many nodes as possible.

## 3. Performance Evaluation

### 3.1 Experiment Set Up

We used three kinds of data sets which were collected from real Internet. The first one is the “King” data that contains measured RTT between any two nodes of Internet 1740 DNS servers using the King method [8]; the second one is a data of measured RTT between 226 nodes of PlanetLab [9]; the third one is “Meridian” data set including measured RTT of 2500 nodes[14].

We developed the PIC simulator based on the description of [3]. In the simulator, nodes were mapped into a 7-dementional Euclidean space, each node had 16 landmarks, of which 4 landmarks were the closest ones, and Simplex Downhill [10] was exploited as the optimization algorithm. For the security policy, the threshold of the relative error mentioned in 2.1 was set to 5% and the repeat time of security policy was set to 5 times. When the security policy is used in PIC, we call it security on; otherwise, it is called security off.

We applied four types of attack: Random Attack, Fixed Point Attack, Colluding Isolate Attack and Combined Attack. In each attack, we repeated 20 times to carry out our experiment by randomly choosing 0%, 10%, 20%, 30%, 40%, 50%, 60% and 70% nodes from all nodes as malicious nodes except the first 16 joined nodes called basic landmarks. Specifically, each attack begins when the number of joined nodes is larger than the required number of basic landmarks.

## 3.2. Performance Metric

In order to understand the condition of TIV in different networks, we analyze the proportion of TIV edge by using the definition of triangulation ratio T=*d*(A,C)/(*d*(A,B)+*d*(B,C)) and utilize equation (4) in section 2.2 to present the severity of TIV.

Then, we use the mean of average relative error as the performance metric in security policy evaluation. In each computation, average relative error can reveal the accuracy of coordinate computation of all the nodes in the system. However, there are many random factors to influent the accuracy so that the average relative error of each computation fluctuates. In order to better evaluate the performance of the system, we run the computation for several times and compute the mean of average relative error to eliminate random factors. The smaller the mean of average relative error is, the more accurate computation is. Average relative error is computed using the first equation as follow and the mean of it uses the second one.

$$\overline{e} = \frac{\left(\sum_{i} \sum_{j(i \neq j)} \frac{d_{i,j}^{m} - d_{i,j}^{p}}{d_{i,j}^{m}}\right)}{M} \tag{5}$$

$$E = E(\overline{e}) = \frac{\sum_{i=1}^{N} \overline{e}_i}{N} \tag{6}$$

## 3.3 Experimental Result

### 3.1.1. Analysis of TIV in different data sets

Based on triangulation ratio definition, we compute how many edges violate triangle inequality in the three data sets respectively. From the result in Table 1, it is obvious that PlanetLab Data and King Data have a similar percentage of TIV edge while Meridian Data includes almost two times TIV edges. That is, the proportion of TIV in Meridian is much larger than the other two.

Table 1. Proportion of TIV edge in Three Data Sets

| Data Set | TIV Proportion |
|---|---|
| PlanetLab Data | 4.4% |
| King Data | 4.11% |
| Meridian Data | 7.84% |

Through the equation (4), we plot a CDF of TIV severity of the three data sets in Fig.1. Considering the meaning of TIV severity metric, we can easily understand that among the three data sets, King data and PlanetLab Data have less severity of TIV than Meridian data set. That means the violation in Meridian is the largest one. Specifically, for 90% edges of the three data sets, the TIV of severity in Meridian is about 0.25 meanwhile that of King and PlanetLab Data are similar, which are only 0.1. Considering results in both Table 1 and Fig.1, Meridian has most TIVs in both quantity and severity. This result will be mentioned from time to time in following discussion.

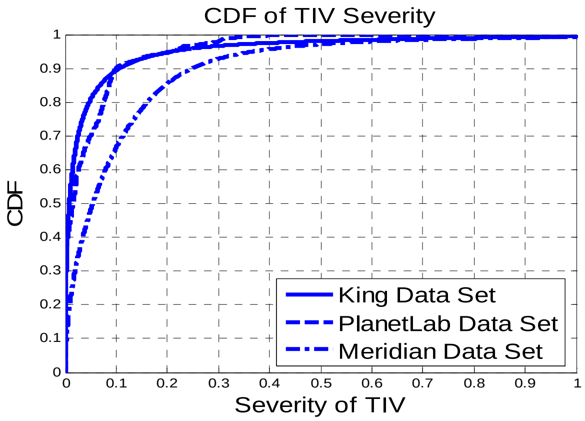

Fig.1. CDF of TIV Severity

### 3.1.2. Random Attack

Random Attack is a simple attack implemented by malicious nodes independently, which can be classified into disorder attack mentioned in 2.3. In this attack, malicious nodes have no special aims except causing high computation error in the system. When malicious nodes are chosen as landmarks by other nodes, they will generate random coordinates and inflate measured distances between them and their victims. In specific, every dimension of random coordinates is in [-250,250] and the inflated distances are 1.5 times more than the true ones.

Fig. 2, 3 and 4 show the mean of average relative error on PlanetLab data, King data as well as Meridian when there are different percentage of malicious nodes in PIC. From them, we observe that when the number of malicious nodes increases, the mean of average relative error rises, which is consistent with our intuition. Especially, when malicious nodes are more than 20%, the error is much larger than that of PIC without malicious node. In other word, Random Attack of malicious nodes would result in high computation errors.

Comparing PIC with security and without security, we find that the security policy can protect PIC from Random Attack. However, when more than 40% malicious nodes present in the system, the difference of the two curves gets smaller with the increase of malicious nodes. Especially in Fig. 3 and 4, when there are more than 50% malicious nodes in the system, the accuracy of PIC with security is almost the same as that of PIC without security. The probable reason for this is that although malicious nodes only attack the nodes choosing them as landmarks, the high computation errors caused by their attack can propagate throughout the whole system. When the number of malicious nodes is less than 40%, their action is distinguished from other nodes and the security policy can detect them easily. So the difference between PIC with security and without security is obvious. However, the increase of malicious nodes leads to more chaos in the system and causes more honest nodes' coordinates inaccurate, which results in that the failure rate in distinguishing malicious and honest ones increases. Thus, when the number of malicious nodes increases in the system, the security policy becomes less effective.

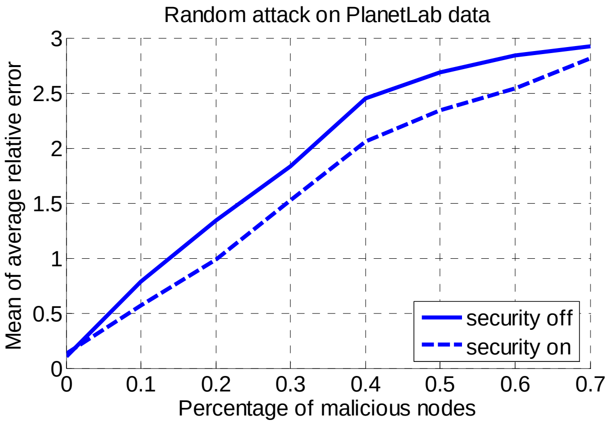

Fig.2. Mean of Average Relative Error of Random Attack on PlanetLab data

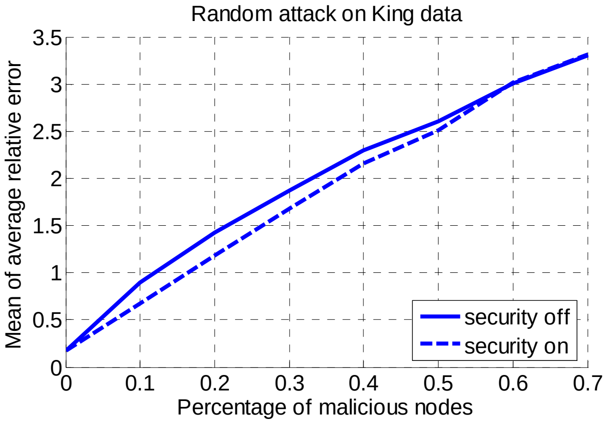

Fig.3. Mean of Average Relative Error of Random Attack on King data

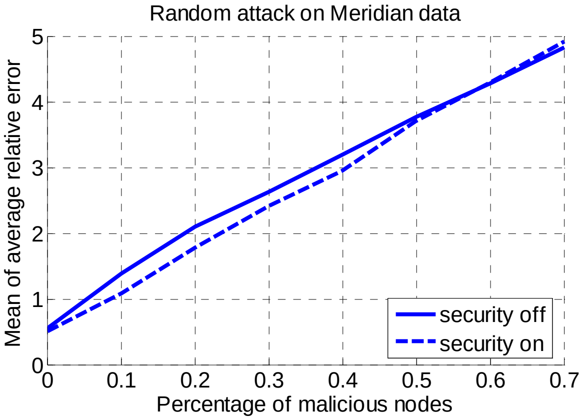

Fig.4. Mean of Average Relative Error of Random Attack on Meridian data

In Fig.2, 3 and 4, it is obvious that the mean of average relative error of Meridian is larger than the other two data sets. For example, without malicious nodes, the error in Meridian is more than 0.5 which is larger than that in King and PlanetLab data sets; while there are 10% nodes conducting random attack, the error increases to more than 1 that is the error of 20% malicious nodes in the other data sets. This is because the TIV severity of Meridian is more and larger than that of King and PlanetLab, thus this will cause high relative error of PIC in Meridian.

### 3.1.3. Fixed Point Attack

In this attack, malicious nodes consult a fixed point as their coordinates before attack and they don't communicate with each other after joining the system. If they are selected as landmarks by other nodes, they will inform the nodes with the fixed coordinates and the real distances between them. In our experiment, we set each dimension of the fixed point to 0.

Fig. 5 presents the results of Fixed Point Attack on PlanetLab data, Fig. 6 depicts the results of King data and Fig.7 is the average relative error of Meridian. Comparing the mean of average relative error when there are malicious nodes in PIC with that of PIC without malicious nodes, we find that Fixed Point Attack causes higher error in coordinate computation. Meanwhile, the error increases with the growth of malicious nodes. In Fig.5, when more than 40% malicious nodes attack PIC, the mean of average error gets lager than 1.

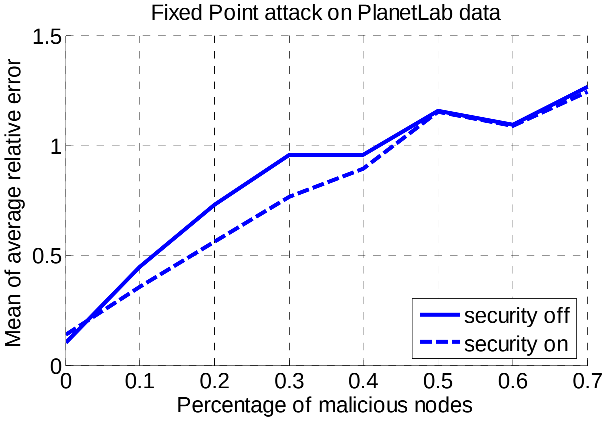

Fig.5. Mean of Average Relative Error of Fixed Point Attack on PlanetLab data

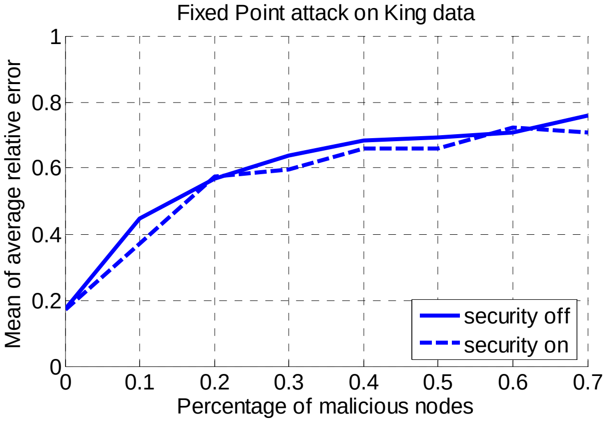

Fig.6. Mean of Average Relative Error of Fixed Point Attack on King data

In Fig. 5, when there are more than 50% malicious nodes, the two curves are almost the same. In Fig. 6, the performance of PIC with security is barely better than without security. Even when the number of malicious nodes is 60%, the performance of PIC with security is a little worse than without security. Security policy in PIC performs worst in Meridian Data Set according to Fig.7, because when there are only 20% and 30% malicious nodes in the system, the accuracy of PIC with security is smaller than PIC without security and when malicious nodes are more than 40%, the security policy works almost the same as PIC with no protection. These figures show the security policy cannot protect PIC well.

Because of more TIVs in Meridian than other two data sets, the relative error is higher even when there are no malicious nodes. And since the security policy aims to detect the nodes which mostly violate triangle inequality and there is higher TIV severity in Meridian, the security policy may falsely remove a severe TIV node as a malicious node from the system, which will cause high inaccuracy. That is probable why when there are less malicious nodes, such as 20% and 30%, the security policy performs even worse than PIC of security off.

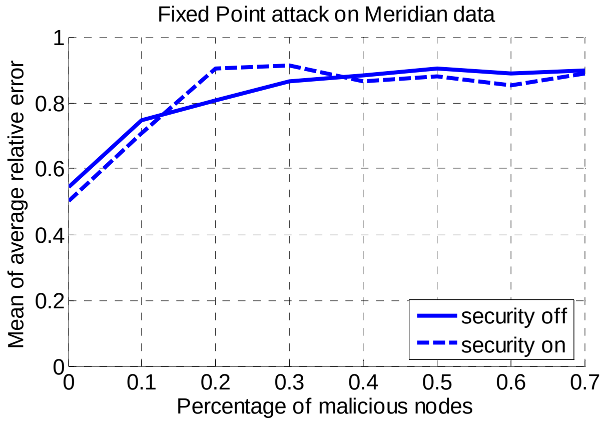

Fig.7. Mean of Average Relative Error of Fixed Point Attack on Meridian data

### 3.1.4. Colluding Isolate Attack

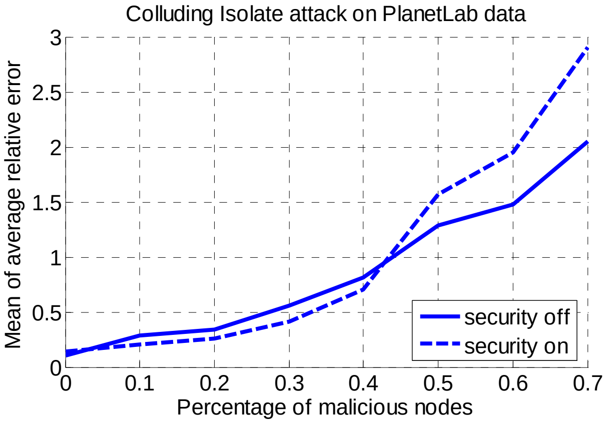

Fig.8. Mean of Average Relative Error of Colluding Isolate Attack on PlanetLab data

In [6], there are two methods to realize Colluding Isolate Attack: one is that malicious nodes move all honest nodes away from the target nodes; the other one is that malicious nodes attract targets into a specific area. We utilize the latter one in this paper: malicious nodes try to inveigle the victims to a settled point in a remote area, each dimension of which is 230. In real network, nodes can delay probes to augment measured distances. Thus, when a node requests information from a malicious node, the malicious node would compute the predicted distance between it and the settled point by using the coordinates. If the predicted distance is less than the real distance between the malicious node and its victim, the malicious node will compute new coordinates to make the predicted distance larger than the real one. Otherwise, it will compute coordinates to get close to the settled point but keep the predicted distance larger than the real one. In both computations, the coordinates change along the line between the malicious node and the settled point. Finally, the malicious node would send the victim its elaborately computed coordinates and the predicted distance between it and the settled point.

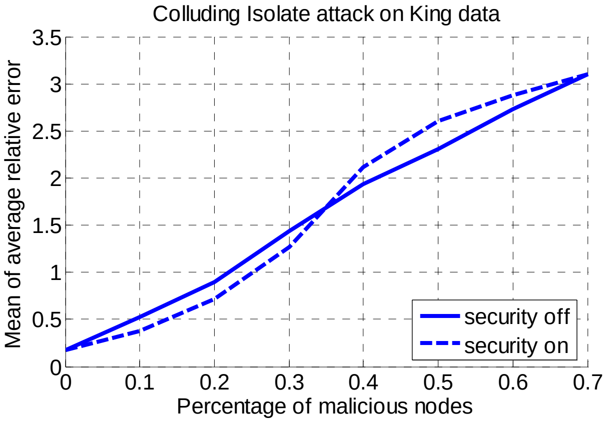

Fig.9. Mean of Average Relative Error of Colluding Isolate Attack on King data

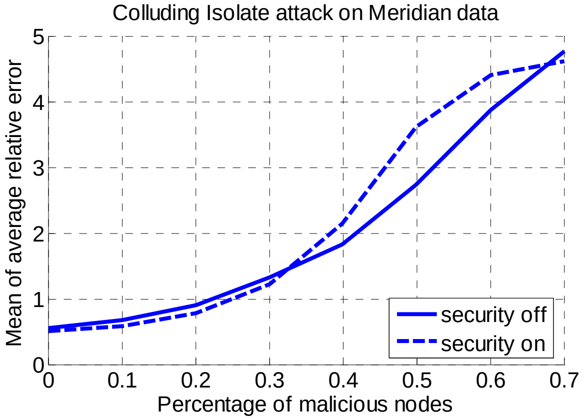

Fig.10. Mean of Average Relative Error of Colluding Isolate Attack on Meridian data

Fig.8, 9 and 10 show that the mean of average relative error rises when the number of malicious nodes increases, which is similar with the results of Random Attack and Fixed Point Attack. We also observe that Colluding Isolate Attack is more powerful in causing high computation errors. Especially, when there are more than 40% malicious nodes in PIC on

PlanetLab data and King data, the error is higher than 1, which is at least 5 times more than that of PIC without malicious nodes. While in Meridian data, the relative error under 40% malicious nodes attack is even higher, about 2.That is, PIC is vulnerable by Colluding Isolate Attack.

Meanwhile, the significant results show that on PlanetLab data, King data or Meridian data, Colluding Isolate Attack can confuse the security policy successfully because this attack attempts to accord with the triangle inequality to avoid being detected. From Fig.8, 9 and 10, we find that PIC with security performs even worse than PIC without security when there are more than 40% malicious nodes in the system. Indeed, when less than 40% malicious nodes attend in PIC, although the security policy can protect the system from the attack, the difference between the two curves in each figure is less than 0.2, which is rather small. When malicious nodes are more than 40%, the capability of the security policy becomes worse---Colluding Isolate Attack on PIC with security results in higher error than without security. Thus, Colluding Isolate Attack is not only powerful in leading to high computation error in PIC but also efficacious to defeat the security policy based on triangle inequality.

Similar as the above analysis, more TIV severity causes higher relative error. In addition, because edges in Meridian violate triangle inequality more than others and colluding attack tends to adapt triangle inequality, severer TIV nodes can delude the security policy of PIC proved by that the maximum difference of the two curves when there are less than 40% malicious nodes in Fig. 10, which is 0.12, is smaller than that in Fig.8 and Fig. 9.

### 3.1.5. Combined Attack

Previously, we study the effect of three attacks on PIC with different percentage of malicious nodes. In Internet, since there are many hackers to attack the system with different intentions, there would be more than one type of attack. The most practical attack should be consisted of different attacks, which we call Combined Attack. In our experiment, we combine three attacks above into this attack and fairly separate malicious nodes into three parts to carry out different attacks. The parameters of each attack are the same as above.

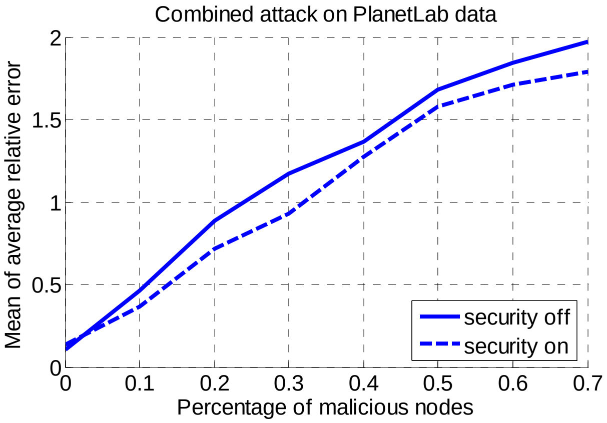

Fig.11. Mean of Average Relative Error of Combined Attack on PlanetLab data

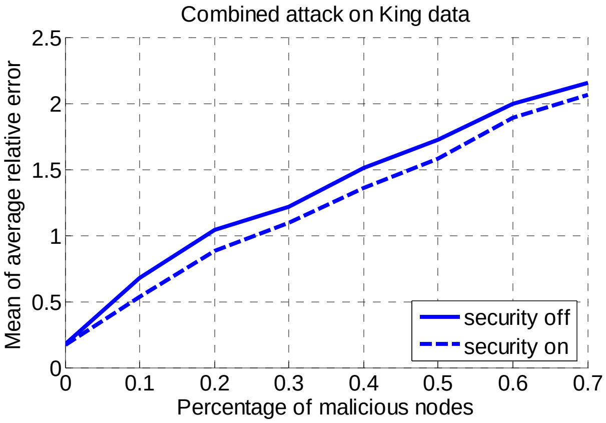

Fig.12. Mean of Average Relative Error of Combined Attack on King data

Fig. 11, 12 and 13 show that the effect of Combined Attack is impactful. When malicious nodes are more than 20%, the mean of average relative error is higher than 1. Moreover, we find that the error grows almost linearly with the increase of malicious nodes. That is, the more malicious nodes attack PIC, the higher error is produced. Combined Attack reveals that in network when different attacks attend in PIC at the same time, the accuracy of coordinat computation in PIC would decline obviously.

What is more, as discussed before, under Combined Attack, higher relative error results from higher TIV severity. The relative error of Meridian data is higher because Meridian has more violation of triangle inequality.

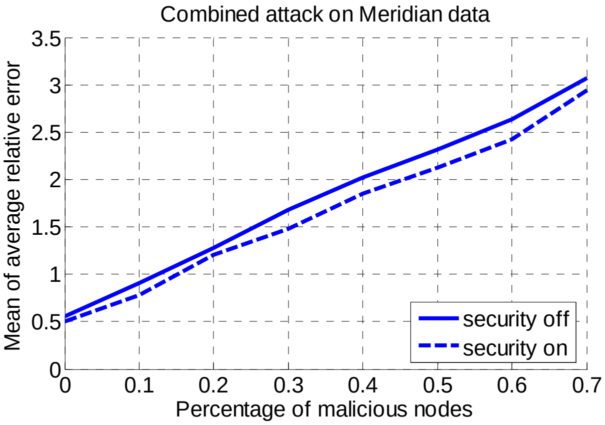

Fig.13. Mean of Average Relative Error of Combined Attack on Meridian data

## 4. Conclusion

PIC, a representative NC system, has designed a security policy based on triangle inequality. In this paper, we study how well PIC security policy can protect PIC by applying four classical attacks on PIC. Considering TIV exists in Internet, we research into the security policy in various network conditions with three different network delay data sets. Analyzing our extensive experimental results, we conclude the following three conclusions.

Firstly, comparing to PIC computed in honest nodes, it is proved that the four attacks will cause higher computation errors in PIC. That is, PIC is vulnerable to attacks. Furthermore, the more malicious nodes exist in network, the higher relative error will be resulted in. We believe that hackers will try to control more nodes to achieve effective attack.

Secondly, according to the analysis, it is concluded that the more violation of triangle inequality exist in network, higher errors of NC computation will be caused no matter PIC with security or not.

Last but not least, comparing with PIC without security policy, the results illustrate that security policy works better under less than 40% malicious nodes, but more than 40% malicious nodes will deteriorate the performance of PIC with security policy. In Colluding Isolate attack, security policy cannot protect PIC at all with more than 40% malicious nodes.

In the future, considering the above conclusions, we will design a more effective security policy to defend NC system from attacks and implement our security policy on NC system in real network to validate its performance.

## 5. Acknowledgement

This work is supported by the National Science Foundation of China (No.60850003, No.60473087, No.60703052) and the National Basic Research Program of China (No.2007CB310806). This paper was presented in part at the 2008 International Symposium on Computer Science and its Applications (CSA'08), Hobart, Australia, Oct.2008.

## 10. References

[1] T. S. Eugene Ng and Hui Zhang. Predicting internet network distance with coordinates-based approaches. In Proceedings of the IEEE INFOCOM, June 2002.

[2] T. S. Eugene Ng and Hui Zhang. A Network Positioning System for the Internet. In Proceedings of the USENIX annual technical conference, June 2004.

[3] M. Costa, M. Castro, A. Rowstron and P. Key. Practical Internet coordinates for distance estimation. In Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS), March 2004.

[4] F. Dabek, R. Cox, F. Kaashoek and R. Morris. Vivaldi: A decentralized network coordinate system. In Proceedings of the ACM SIGCOMM, August 2004.

[5] Y. Chen, Y.Q. Xiong, X.H Shi, B.X. Deng and X. Li. Pharos: A Decentralized and Hierarchical Network Coordinate System for Internet Distance Prediction. In Proceeding of the 50th Annual IEEE Global Telecommunications Conference (GLOBECOM), November 2007.

[6] M. A. Kaafar, L. Mathy, T. Turletti and W. Dabbous. Real attacks on virtual networks: Vivaldi out of tune. In Proceedings of the SIGCOMM workshop on Large Scale Attack Defense (LSAD), September 2006.

[7] A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In Proceedings of IFIP/ACM International Conference on Distributed Systems Platforms, November, 2001.

[8] K. P. Gummadi, S. Saroiu, and S. D. Gribble. King: Estimating Latency between Arbitrary Internet End Hosts. In Proceedings of SIGCOMM Internet Mesasurement Workshop (IMW), November 2002.

[9] PlanetLab: http://www.planet-lab.org/

[10] J. A. Nelder and R. Mead. A simplex method for function minimization. Computer Journal, 7:308–313, 1965.

[11] M. A. Kaafar, L. Mathy, T. Turletti and W. Dabbous. Virtual networks under attack: Disrupting internet coordinate systems. In Proceedings of Second CoNext Conference, 2006.

[12] H. Zheng, E. K. Lua, M. Pias, and T. Griffin. Internet routing policies and round-trip times. In the 6th anual Passive and Active Measurement Workshop, 2005.

[13] G. Wang, B. Zhang, and E. Ng. Towards network triangle inequality violations aware distributed systems. In Proceedings of the ACM/USENIX Internet Measurement Conference (IMC'07), Oct 2007.

[14] B. Wong, A. Slivkins, and E. G. Sirer. Meridian: A lightweight network location service without virtual coordinates. In *Proceedings of ACM SIGCOMM*, August 2005.