

Certified Internet Coordinates

Mohamed Ali Kaafar
INRIA Sophia Antipolis, FR
mkaafar@sophia.inria.fr

Kave Salamatian
LIP6, FR and EPFL, CH
kave.salamatian@lip6.fr

Laurent Mathy^{*}
Lancaster University, UK
laurent@comp.lancs.ac.uk

Thierry Turletti
INRIA Sophia Antipolis, FR
turletti@sophia.inria.fr

Chadi Barakat
INRIA Sophia Antipolis, FR
barakat@sophia.inria.fr

Walid Dabbous
INRIA Sophia Antipolis, FR
dabbous@sophia.inria.fr

ABSTRACT

We address the issue of asserting the accuracy of Internet coordinates advertised by nodes of Internet coordinate systems during distance estimations. Indeed, some nodes may even lie deliberately about their coordinates to mount various attacks against applications and overlays.

Our proposed method consists in two steps: 1) establish the correctness of a node's claimed coordinate by using the Surveyor infrastructure and malicious embedding neighbor detection proposed in our previous work on securing the coordinates embedding phase; and 2) issue a time limited validity certificate for each verified coordinate.

Validity periods are computed based on an analysis of coordinate inter-shift times observed by Surveyors. By doing this, each surveyor can estimate the time until the next shift and thus, can limit the validity of the certificate it issues to regular nodes for their calculated coordinates. Our method is illustrated on a trace collected from a Vivaldi system deployed on PlanetLab, where inter-shift times are shown to follow long-tail distribution (lognormal distribution in most cases, or Weibull distribution otherwise).

We validate the effectiveness of our method by measuring the impact of a variety of attacks on distance estimates.

1. INTRODUCTION

Internet coordinate systems (e.g. [1, 2, 3, 4, 5, 6, 7]) have been proposed to allow for distance (Round-Trip Time, shortly RTT) estimation between nodes, in order to reduce the measurement overhead of many applications and overlay networks. Indeed, by embedding the Internet delay space into a metric space – an operation that only requires each node in the system to measure delays to a small set of other nodes (its neighbors), nodes are attributed coordinates that can then be used to estimate the RTT between any two nodes, without further measurements, simply by applying the distance function associated with the chosen metric space to the nodes' coordinates.

The security of the embedding phase, or in other words the

coordinate computation phase, of Internet coordinate systems has received some attention in the form of either simple mechanisms built into embedding protocols [3, 4] or a more general, embedding protocol-independent cheat detection mechanism based on a Surveyor infrastructure [8]. All of these mechanisms use distance measurements between two nodes to, one way or another, assess the plausibility of the corresponding distance estimation based on the nodes' current coordinates. In cases where the discrepancy between the measured and estimated distances is deemed too important, the node carrying out the test removes its correspondent node from its set of neighbors. By doing this, it avoids adjusting its coordinate in response to potentially malicious information.

However, ultimately, Internet coordinate systems are used to estimate distances between nodes, based on their coordinates only, even and all the more so if these nodes have never exchanged a distance measurement probe. Whatever mechanism is used to obtain a node's coordinate (e.g. direct exchange, DNS-like repository, etc.), each node must somehow report its own coordinate computed during the embedding phase of the system. This is because, for scalability reasons, nodes compute their own coordinates in a distributed way. This, of course, provides a malicious node with an opportunity to strike: in order to achieve some application-dependent goal or advantage (e.g. free-riding, denial-of-service, isolation, ubiquity gift, etc), a node can repeatedly lie about its coordinate. Simply lying about its coordinate could seriously disrupt the operations of Internet applications relying on coordinate-based systems for distance estimation. Several studies have quantified the impact of cheating on topology-aware Internet applications, and have shown that simple attack strategies can prove to be very effective [9, 10, 11]. This paper addresses the question of guaranteeing the veracity of the coordinates advertised by nodes.

To do so, we propose to leverage the Surveyor infrastructure and embedding cheat detection test proposed in [8]. More precisely, we propose that several Surveyors measure their distance to a node in order to verify the correctness of its claimed coordinate (using the cheat detection test from [8]). If all Surveyors agree that this coordinate is the node's true

^{*}L. Mathy was also partly supported by the University of Liège, Belgium.

coordinate, a time limited validity certificate, including the certified coordinate, is issued to the node.

A certificate validity time period is necessary because, due to dynamic network conditions, nodes’ coordinates vary in time. Upon a coordinate change, an honest node would stop using its current certificate and seek a certification of its new coordinate. On the other hand, a malicious node could keep using a certificate related to a previous position, hence a careful balance between scalability and certificate validity is desirable. To achieve this, one of our contributions is to study the coordinate inter-shift time (i.e. the time between coordinate changes at a node) as observed for a Vivaldi system running on PlanetLab. We found that the coordinate inter-shift times at most nodes follow a lognormal distribution, with the rare cases when this distribution is inappropriate being accounted for by a Weibull distribution (note these are both long-tail distributions).

The remainder of the paper is as follows: in section 2, we study and characterize the coordinate inter-shift times and show that these times observed at Surveyors can adequately and statistically model inter-shift times at closeby nodes. Section 3 describes the certification procedure in detail, while performance evaluation of our proposal in the context of various attacks is presented in section 4; section 5 concludes the paper.

2. COORDINATES EVOLUTION MODEL

Our goal is to characterize and ascertain forward validity of nodes’ coordinates during the distance estimation phase. We therefore concentrate on tracking the evolution in time of coordinates along the different axis of the coordinate space by observing their evolution in a clean system (without any malicious activity).

2.1 Experimental Set-up

The measurement used in this study comes from traces obtained by running the Vivaldi System on 600 planetLab machines through a period of 18-days. These traces were then acquired in a clean environment with no malicious node. We also used live-PlanetLab experiments, as a Vivaldi service deployed over a three-weeks period, to demonstrate the validity and effectiveness of our proposal to deal with various attacks.

In our initial trace, we kept the coordinates of 450 nodes, as several nodes, that have been down for more than one week or experienced connectivity troubles, have been filtered out. Some of them even do not respond to different RPC requests sent by the Vivaldi service.

In this trace, Vivaldi is using a 3-dimensional Euclidean space for the embedding. Each 10 minutes, corresponding to an embedding step, nodes are adjusting their coordinates based on a one-to-one interaction with another node, called a peer node. Finally, Vivaldi nodes are filtering the stream of latency measurements from a remote node and turn these into expected latency values, based on a Moving Percentile

(MP) filter, a variant on the Moving Median filter, used to filter out heavy-tailed errors [12, 13]. This filter acts as a low-pass filter, where anomalies are ignored while a baseline signal passes through. This allows Vivaldi to smoothly adapt to shifts in the baseline (that BGP route changes cause for example). The goal of the latency filter is to summarize the measurements, providing a current and stable description of the expected latency between any two nodes. Ledlie et al. [12] found that using this simple latency filter with the a size of the history window equal to four observations and a percentile of 25% was a good predictor for the next observed latency (i.e. consider the minimum value of a four-latency measurements stream, for coordinates computation).

The PlanetLab experiments were conducted over a set of 280 PlanetLab nodes spread world-wide, running Vivaldi, as a coordinate-embedding service. For the purpose of our experimentations, we slightly modified the logging functions of the Vivaldi protocol. Each node is running several instantiations to allow us experimenting different parameters in similar conditions. Nodes are then updating their coordinates as needed, depending on the embedding step defined in each instantiation of the Vivaldi protocol. In the same way, the dimension of coordinates varies from one instantiation to another on the same node ¹, and a node is acting as a malicious node or as a honest ‘normal’ node.

Each node had 20 neighbors (i.e. was attached to 20 springs), 10 of which being chosen to be closer than 50 ms. The constant fraction C_c for the adaptive timestep (see [5] for details) is set to 0.25. When needed, Surveyor nodes were chosen that represent 8% of the overall population [8].

2.2 Observations

Figure 1 shows a typical evolution of the coordinates of a Vivaldi node. Each sub-figure depicts the evolution of one coordinate component along the three axis.




Figure 1: Typical Variations of a node’s coordinate in the Vivaldi System.

¹We experimented with different Euclidean embedding spaces and a 2-dimensional coordinate space augmented with a height vector.

We observe that the system after roughly 250 embedding steps, reaches a stationary regime, but coordinates continue to fluctuate. Looking further in the entire set of nodes' coordinates, we observe that coordinates steadily variate, moving away from the origin (0^3) of the coordinate system. The rate of these variations is low enough to allow distance estimation at the application level, but this prevents us from certifying absolute coordinates.

More specifically, because at any instant in time, the RTT that can be measured between two nodes depends on the state of the network (e.g. traffic load, state of queues in routers, etc), the exact value of the RTT varies continuously. However, it has been shown that RTT values in the Internet exhibit some stability in a statistical sense [14], with the statistical properties of RTTs exhibiting no significant change at timescales of several minutes. It is that property that embedding systems exploit to provide good distance estimates while only needing to have nodes adjust (recalculate) their coordinates on a periodic basis. Consequently, the node's coordinate can be viewed as a discrete stochastic process, embedded at the instants of updates.

Regardless of the dimensionality used by the coordinate-systems, our main goal is to assign to any coordinate given by a node, a reliability value that is telling the likelihood that this coordinate is still valid and has not changed. For this purpose we observe for our set of 450 planetlab nodes the inter-shift time distribution, corresponding to the amount of time (in terms of embedding steps intervals) during which, nodes stick to their positions, i.e. the coordinates do not change. This distribution is denoted T_i for each node i . It is important to note that although we observed that in our traces, a variation of one coordinate component was synonym to the variation of both others, we consider the inter-shift time as the laps of time corresponding to the non variation of all the coordinate components of this node.

Basically, we would like to determine which probability distribution is suitable to describe the inter-shift times. In the following, we will use our empirical data sets of inter-shift times to find the probability distribution that best describes the distribution values of each T_i .

2.3 Inter-shift Time Distribution Fitting

For choosing the best suitable distribution, we use a set of candidate distributions containing lognormal, Weibull, Rayleigh and Gamma distributions².


For each one of the 450 nodes in the dataset, we apply a two-step procedure. In the first step, we derive the maximum likelihood estimates of parameters of each distribution in the candidate set. The likelihood is the probability to get the observed inter-shift times for some hypothetic distribution. The estimates of the distribution parameters are then the values that maximize their likelihood.

In a second step, we used goodness of fit tests to evalu-

²The Gaussian distribution was not tested because the empirical distribution was not symmetrical around a mean.

ate if the hypothesis that the observed values T_i come from the candidate distribution can be rejected or not. The goodness of fit evaluation was done using the popular and robust Kolmogorov-Smirnov (K-S) test [15], applied at a significance level of 5%.

We show in figure 2, the results of the fitting done on one



(a) Histogram of the data and (b) Empirical CDF and fitted PDF of fitted distributions CDFs

Figure 2: Density plots and CDF of a typical inter-shift distribution (Node 181).

of the datasets relative to node 181. Simple inspection by eye shows that the lognormal distribution produces the best fits among the set of distributions. This is validated by the K-S test that rejects all the candidate distributions but the lognormal.

Using the fitting procedure described above, we tried to fit the inter-shift datasets. The first interesting result we found is that all of the empirical distributions examined can be fitted to a known distribution. A large majority of distributions can be fitted into a lognormal distribution. The lognormal hypothesis was rejected for only 5 datasets out of the 450. Looking further in these 5 inter-shift datasets, we observed that they have a good fit with the Weibull distribution. Table 1 gives a summary of our findings for the Kolmogorov Smirnov goodness of fit tests.

2.4 Correlation between Surveyors and Nodes inter-shift Distributions

Having shown that inter-shift time distribution of our population of nodes in the Vivaldi system is compatible with either a lognormal distribution in most cases or a Weibull distribution otherwise whose parameters can be obtained by a maximum likelihood method, the next question is how well the shifts as observed by Surveyor nodes can be used as representative of the shifts of regular nodes. Basically, if the inter-shift distribution of a node as seen by the surveyors is the same as the real inter-shift distribution, the former may be used as a reference to validate the node's coordinate.

The verification of this hypothesis is done by comparing the sequence of inter-shift as seen by surveyors and the real inter-shift of a node and asking if the hypothesis that these two sequences come from the same distribution can be rejected or not. The latter test is done using a two-sample Kolmogorov Smirnov Test (with a significance level of 5%) that precisely gives an answer to the previous question. For

Table 1: Results using the Kologorov Smirnov Test for fitting the inter-shift Times data.

Fitted Distributions	% of samples that passed the test
<i>log-Normal</i>	445/450
<i>Rayleigh</i>	2/450
<i>Weibull</i>	5/450
<i>Gamma</i>	6/450

each of the 35 surveyor nodes, we applied therefore 415 two-sample K-S test.

To answer this question, we analyzed for each surveyor nodes the likelihood that the two-samples KS-test is rejected as a function of the distance (measured as an RTT) between the surveyor node and the tested regular node. The likelihood is obtained as the ratio between the number of nodes with a given RTT that reject the test and the overall number of regular nodes. Figure 3 shows this rejection ratio vs. the distance (measured as an RTT) between a node and the corresponding Surveyor, as observed during the PlanetLab experiment.




Figure 3: Correlation between 'Nodes-Surveyors' RTTs and the rejection Ratio.

Intuitively, a Surveyor should have tendency to have the same inter-shift distribution as nodes that are close by in terms of RTT, as they are more likely to experience similar dynamics of the coordinate system. Figure 3 validates this intuition and shows that better locality between a node and its Surveyor yields more accurate fittings.

We therefore apply the following heuristics: the inter-shift time distribution of the closest Surveyor is used as the representative distribution for regular nodes.

3. COORDINATE CERTIFICATION

The method we propose to certify Internet coordinates consists in two steps:

1. the node coordinate verification test;
2. computation of an estimated validity period for this coordinate.

The coordinate verification test leverages the Surveyor infrastructure and malicious embedding neighbor detection proposed in [8], while the validity period estimation is based on the results presented in section 2.

3.1 Coordinate Verification

3.1.1 Principle

The method for malicious embedding neighbor detection proposed in [8] is based on a model of the evolution of a nodes's observed relative errors, during embedding steps. These relative errors, defined as $\frac{||X_i - X_j|| - RTT_{ij}}{RTT_{ij}}$, where X_i and X_j are the nodes' coordinates and RTT_{ij} the measured distance between the nodes, have been shown to be trackable by a Kalman filter in the case of a system without malicious nodes. The idea is thus to have a set of trusted nodes (a.k.a Surveyors), well positioned in the network and acting as vantage points, that compute their coordinates by embedding with each other exclusively and thus observe the evolution of their relative errors in a clean system without influence from malicious nodes. It was also shown that the Kalman filters calibrated on such nodes can advantageously track the relative error model at nearby untrusted nodes. Each Kalman filter is capable of providing a predicted relative error (based on the history of observed measured errors) as well as the distribution of its innovation process which is the deviation between the measured and predicted relative errors expected at each embedding step. The basis for the malicious embedding neighbor detection test is then a simple hypothesis test on this innovation process (see [8] for details).

It is important to note that this malicious embedding neighbor detection test is based on relative errors, and thus simply tests the consistency between the estimated and measured distances between two nodes. In particular, this test is not able to evaluate the truthfulness of a node's coordinate. Indeed, if during an embedding step a node fakes its coordinate but at the same time delays the measurement probes in a way consistent with its faked position in the coordinate space (based on the knowledge of the correspondent's coordinate, as well as its own true and fake coordinates), then the resulting relative error measured between the two nodes is reduced (see figure 4). Indeed, if a node at coordinate X_i fakes a coordinate X'_i , then its real distance to another node (e.g. a

Surveyor) at coordinate X_j is $RTT_{i,j}$ (this RTT is the time the probe actually travels between the 2 nodes). To be consistent with its faked position, the faking node should delay the measurement probe by $\|X'_i - X_j\| - \|X_i - X_j\|$, so that the measured RTT, $RTT'_{i,j} = RTT_{i,j} + \|X'_i - X_j\| - \|X_i - X_j\|$. The relative error that should have been measured is clearly $\frac{\|X_i - X_j\| - RTT_{i,j}}{RTT_{i,j}}$, while the relative error measured due to the faked position is $\frac{\|X'_i - X_j\| - RTT'_{i,j}}{RTT'_{i,j}}$. Simple substitution allows the measured relative error to be rewritten as $\frac{\|X_i - X_j\| - RTT_{i,j}}{RTT'_{i,j}}$, which is clearly lower than the relative error that would be measured if the node did not fake its coordinate (since $RTT'_{i,j} > RTT_{i,j}$). Therefore, a node faking a position further away from the testing node will never fail a test that it wouldn't have failed in the first place if it wasn't faking.




Figure 4: Fake Coordinate and Consistent Relative Error

Consequently, to verify a node's coordinate, several such tests must be performed from vantage points (Surveyors) *surrounding* the node (see figure 6). In this case, a node could easily fake its coordinate and consistently delay probes so that it moved away from some Surveyors without being noticed. But such fake position would necessarily also result in the node moving closer to some other Surveyors and failing the corresponding malicious embedding neighbor tests as it is next to impossible to “speed up” a distance probe protected by the simplest of mechanisms (e.g. hashing, simple encryption, random probe numbers, etc). A node must be surrounded by at least one more Surveyors than there are dimensions in the space.

If the malicious embedding neighbor test is negative at each Surveyor chosen to surround the node (i.e. the relative error observed between the Surveyor and the node is considered normal), then the coordinate claimed by the node is considered correct. Note that this test is different from a normal “triangulation” approach, where the measured distances between the node and the Surveyors would be used alongside the Surveyors' coordinates to determine the node's own coordinate. Indeed, our test is in fact made up of multiple, independent tests on the plausibility of the observed relative

errors and provides our method with an important resistance to the triangular inequality violations (TIVs) [16, 17] that can be commonly encountered in the Internet. This is because the Kalman filters underlying the tests are calibrated during normal embedding of the Surveyors, and thus in conditions where TIVs are naturally encountered, so the system noise resulting from these TIVs is therefore implicitly taken into account in the relative error tracking. We do not claim that our test is immune to the problems caused by TIVs (and these TIVs will be responsible for some of the false positives of our test), but it is nevertheless much less sensitive to them than a geometric approach like triangulation would be.

If a node cannot be surrounded by Surveyors (i.e. there is at least one axis along which the corresponding coordinate component of the node is either greater or smaller than the corresponding coordinate component of all Surveyors), then the node's coordinate cannot be verified. This is because the node could easily fake its position by moving away from the Surveyors along this axis, without being detected, by using the technique described in figure 4 with each Surveyor. It is therefore important to have Surveyors placed near the coordinate space boundaries, as well as inside the space (otherwise part of the nodes population may never be able to get coordinate certificates). More precisely, the convex hull³ of all nodes (including the Surveyors) should be composed of Surveyors only. Or seen differently, any node lying outside the convex hull of the set of Surveyors can never be surrounded and thus can never be issued a coordinate certificate. For this paper, and to reflect a plausible deployment scenario, some of the Surveyors were deliberately placed near the space boundaries, so that the Surveyors' convex hull encloses most, but not all, nodes, while the other Surveyors were chosen at random (see Figure 5).




Figure 5: Distribution of regular and Surveyor nodes in a 2d space (distances are in ms).

3.1.2 Protocol

³The convex hull of a set of points is the smallest convex set containing all the points. Think of it as being the smallest “enclosure” that can be formed with some of the points and that has no points outside it [18].




Figure 6: Surveyors surrounding a node for coordinates verification

A node who wish to have its coordinate certified contacts the known closest Surveyor to its claimed coordinate. If this Surveyor is not the closest, the node is redirected to the closest one known by the Surveyor. For this, as well as the selection of Surveyors surrounding the coordinate claimed by a node to happen, Surveyors exchange their coordinates using a gossiping protocol.

Based on its knowledge of the position of other Surveyors, as well as on the coordinate of the node to be certified, the certifying Surveyor selects a set of Surveyors (a.k.a surrounding Surveyors) that surround the node’s claimed position (possibly including itself) and informs these of the node’s claimed coordinate so to ensure they all use the same node’s coordinate for their distance estimates during their malicious embedding neighbor detection test.

Recall that Surveyors compute their own coordinate by using each other exclusively as embedding neighbors [8]. This gives the Surveyors the view of a clean system without malicious insider attacks. Therefore, if Surveyors run their own malicious embedding neighbor detection test at each embedding step, all such tests should ideally be negative as Surveyors only have trusted and honest neighbors (other Surveyors). Unfortunately, no test is perfect and some amount of the tests carried out by each Surveyor will wrongly identify the neighbor as malicious. Such occurrence is a false positive and the Surveyor will take no action about the said neighbor (or more precisely the particular embedding step with this neighbor). However, carrying out such tests at every embedding step provides the Surveyors with estimates of two important test statistics: the false positive test ratio (FPTR – i.e. the percentage of the tests that were positive and wrongly identified a fellow Surveyor as malicious) and the true negative test ratio (TNTR – i.e. the percentage of the tests that were negative and thus correctly identified the

fellow Surveyors as honest nodes).

Let Y_i be the indicator random variable that represents the outcome of a malicious embedding neighbor detection test at the i^{th} Surveyor (testing another Surveyor), with:

$$Y_i = \begin{cases} 0 & \text{if the neighbor identified as honest} \\ 1 & \text{if the neighbor identified as malicious} \end{cases}$$

Taking as null hypothesis H_0 that the tested node is honest (which is always the case when Surveyors are tested), the true negative test ratio (TNTR) estimate p_i at the i^{th} Surveyor is $\mathbb{P}\text{rob}\{Y_i = 0|H_0\}$, the number of tests that were correct divided by the overall number of tests carried out. The FPTR is then obviously $1 - p_i$.

After performing the requested malicious embedding neighbor detection test on the node whose coordinate is to be certified, the surrounding Surveyors return the result of their test, along with their estimated TNTR, to the certifying Surveyor. If every test returned is negative (i.e. each surrounding Surveyor considered the node as honest), then the node’s coordinate is deemed correct and verified and the certifying Surveyor proceeds to the second step of the certification described in section 3.2.

On the other hand, if at least one of the surrounding Surveyor returns a positive test, that is, did consider the node as potentially malicious because of too much of a deviation between the measured relative error and the expected one, the certifying Surveyor must decide whether to declare the node’s coordinate as suspicious and thus refuse to issue a certificate, or whether further test should be carried out. To do so, the probability that the node, and its claimed coordinate, have been identified mistakenly as suspicious by the surrounding Surveyors is computed. This probability is simply $1 - \prod_{i \in \xi^j} p_i$, where ξ^j is the set of surrounding Surveyors chosen to verify the claimed coordinates of the node at this round of testing. If the overall probability that the node has been mistakenly classified as suspicious is greater than a given significance value γ , that is if $\prod_{1 \leq j \leq \mathcal{N}} (1 - \prod_{i \in \xi^j} p_i) > \gamma$, where \mathcal{N} is the number of test rounds that have been carried out so far, then the certifying Surveyor starts another test round with a new set of surrounding Surveyors. Note that the sets of selected surrounding Surveyors at each round are not necessarily disjoint, although such property is desirable.

In this paper, we used $\gamma = 1\%$ and limited \mathcal{N} to 6 (i.e. a node is refused a coordinate certificate if the probability of mistaken refusal falls below 1% and/or the node fails 6 consecutive test rounds).

3.1.3 Evaluation

In this section, we seek to evaluate the effectiveness of our proposed coordinate verification method.

We first seek to validate the assumption that the FPTR values measured during embedding at Surveyors provide good estimates for the real FPTR values at these Surveyors. To do so, we let a vivaldi system, without cheat, converge and

run on PlanetLab for over 2500 time ticks (i.e. embedding periods). the Surveyors measured their estimated FPTR by carrying out a malicious embedding neighbour detection test at every embedding step. At the end of the experiment, each Surveyor also measured its real FPTR by running a malicious embedding neighbour detection test to every other nodes in the system. Since this system does not have any malicious node in it, a failed test is a false positive. The CDF of the differences of these 2 values at each Surveyor is shown in figure 7. We see that the difference between the estimated and the real FPTR are mostly within less than 1% of each other, confirming that our proposed estimation method yields reliable FPTR estimates. Even in the cases where the FPTR estimates differ more than the real value, these will only affect the coordinate verification tests in which the corresponding Surveyors take place: in these cases the coordinate verification test will be slightly more aggressive than it ought to (since the FPTR estimate is smaller than the real value), favouring security.





Figure 7: CDF of False positive probability differences.

Next, we seek to further understand the behaviour of false positive and true negative occurrences. Figure 8(a) is a scatter plot of the distance between a node and its Surveyor, whenever the malicious embedding neighbour detection test through a false positive at the Surveyor. This figure clearly indicates that Surveyors hardly ever experience wrong test results when testing nearby honest nodes. Complementarily, Figure 8(b) shows that Surveyors are much more successful at identifying nearby honest nodes correctly. These results indicate that striving to choose surrounding Surveyors as close as possible to the node whose coordinate are being verified will increase the effectiveness of the coordinate verification test (by reducing the occurrences of false positive in the multiple test, through reduction of false positive occurrences in the component tests making this multiple test up). This is therefore the strategy adopted in the rest of this paper.

To assess the overall effectiveness of the proposed coordinate verification method, we then make all nodes try to move



(a) Correlation between False Positives and RTTs (b) Correlation between True Negatives and RTTs

Figure 8: Correlation between distance and test performance.

away in all direction, from their position in the converged system. In this experiment, each node varies the direction of its move one degree at a time. Along each direction, it tries to move a distance drawn at random from the distribution of the innovation process computed by the Kalman filter calibrated at the nearest Surveyor [8], in a bid to outsmart the malicious embedding neighbour detection test. After each move, the node requests a re-certification of its coordinate. The displacement achieved by the node until the corresponding fake coordinate is declared incorrect by our verification method is then recorded. Figure 9 shows, the average, minimum and maximum displacement observed across all nodes, as a function of the direction of movement. Note that this figure only take into account the nodes that are eligible for certification (i.e. the nodes that lie inside the Surveyors convex hull, see section 3.1.1), meaning that 13 nodes are excluded from the experiment.




Figure 9: Undetected displacement as a function of direction.

We see that the maximum displacement achieved by any node is below 10 ms, while the average is around 6 ms. Such deviations are of the order of the distance prediction accuracy (they will not decrease the accuracy of the coordinate system in any significant way) and will be acceptable to many applications.

Finally, we experimented with a simple attack, carried out by a growing malicious node population that has access to

the coordinates of all nodes in the system. The malicious nodes compute the centroid of the overall node population, and then try to move in a random direction away from this centroid (adding 2 seconds) in order to be isolated. Figure 10 shows the detection rate, that is the percentage of certificate requests for faked coordinates that was denied, as a function of the malicious population size, for various dimensions of the coordinate space. Note that although these curves show a slightly decreasing trend, a smaller percentage of an increasing number of requests for faked coordinates does mean, in most cases, an increasing number of denials. With over 95% detection rates in most cases, the coordinate verification test can be considered as highly efficient.




Figure 10: Self isolation: Detection probability.

However, tests may achieve high detection rates by being overly aggressive towards honest nodes that do not cheat. The false positive rate, that is the percentage of certificate requests for real coordinates that were wrongly denied, is a measure of this aggressivity towards non malicious nodes and should be kept as low as possible. Figure 11 shows the measured false positive rate, as a function of the size of the malicious population, for various dimensionality of the coordinate space. Note that these curves depend on the performance of the test only, and not on the activities of the malicious nodes. Also, note that as the population of malicious nodes increases, the corresponding population of honest nodes decreases, so an upward trend actually corresponds to fewer wrongly denied certification requests to honest nodes. With a false positive rate lower than 6% in all cases, our test can be considered as moderately, and acceptably, aggressive.

In light of the evaluation results presented in this section, we conclude that our proposed test for coordinate verification exhibits good performance and is fit for purpose.

3.2 Certificate Validity Computation

After the correctness of a node’s advertised coordinates has been asserted, the next step is to issue the node with a certificate proving that the enclosed coordinate has been verified. This certificate will be delivered by the certifying Surveyor (i.e. usually the Surveyor closest to the node in the coordinate space, see sections 2.4 and 3.1.2).

As a coordinate certificate is associated with a particular




Figure 11: Self Isolation: False Positive Ratio.

position in the coordinate space that the node occupies or has occupied, one could expect that nodes request new certificate on moving to a new position. While this is certainly the behaviour anticipated from honest nodes that are interested in being able to prove their true position, malicious nodes might decide not to do so. Indeed, a malicious node is probably much more interested in faking its true position, and doing so with a certificate “proving” a position which isn’t really its own, whatever this coordinate might be, is probably a bonus for such a node. While the coordinate verification protocol described in section 3.1.2 has been designed to prevent, as much as possible, malicious nodes from acquiring certificates for fake coordinates that they may choose, the problem of nodes using certificates for positions that have meanwhile changed, is still to be addressed.

The obvious way to solve this “stale” certificate problem, is to assign a reliability to the certificate. The reliability of the certificate decreases with time from its issuance time. When it crosses a certain threshold p_{th} the certificate should be invalidated and eventually reissued. There is a trade-off between certificate precision (and therefore security) and frequency of coordinate certification that is controlled by the reliability p_{th} . Using larger value of p_{th} , leads to higher reliability for certificates but at the cost of frequently reissuing certificates that are still valid. On the other hand lower p_{th} results in lower reliability, but also reduces the load on Surveyors who would receive certificate requests less frequently. The issue here is really to find the right trade-off between scalability (by limiting the rate or amount of certificate requests) and security (by limiting the time lapse during which a malicious node may be able to use an old certificate after its coordinate has changed). We will therefore seek to exploit the results on coordinate inter-shift times presented in section 2.3 to compute certificate validity periods.

Note that this section assumes that all nodes in the system are “loosely” time synchronized: since coordinate inter-shift times have been shown to take values that are usually measured in minutes (see section 2), as long as all clocks are synchronized with an accuracy exhibiting a smaller timescale (say a few seconds), the notion of time in the system can then be considered unique. This time synchronization can easily be obtained through using NTP (at least about survey-

ors).

3.2.1 Principle

The problem of computing a validity period for coordinate certificates can be formalized through reliability theory. Let's define the survival function of the coordinate of node i as the probability, $S_{T_0^i}^i(\Delta) = \mathbb{P}\text{rob}\{T^i > T_0^i + \Delta\}$, that the next change of coordinate occurs later than Δ time units after the last coordinate change, happening at time T_0^i . The survival function is thus related to the inter-shift time cumulative distribution $F^i(\Delta) = \mathbb{P}\text{rob}\{T^i \leq T_0^i + \Delta\}$ through $S^i(\Delta) = 1 - F^i(\Delta)$.

Recall from section 2.4, that the inter-shift times observed at a Surveyors are similar to those observed at nearby nodes. Hence, the inter-shift time distribution at a certifying Surveyor is the distribution used to compute the validity of the certificates it issues (since it issues certificates to the nodes that are closest to it than to any other Surveyors).

The survival function can be used to compute the validity of a certificate, that is the time remaining until the next coordinate change. The probability that the next position change occurs at or before time $\tau + \Delta$, given that the certificate is being issued at time τ is just:

$$\begin{aligned} \mathbb{P}\text{rob}\{T < \tau + \Delta | T > \tau\} &= \frac{\mathbb{P}\text{rob}\{\tau < T < \tau + \Delta\}}{\mathbb{P}\text{rob}\{T > \tau\}} \\ &= 1 - \frac{S^i(\tau + \Delta - T_0^i)}{S^i(\tau - T_0^i)} \end{aligned}$$

In fact, we use the above probability, computed at a Surveyor whose survival function and last coordinate change time are $S^i(\Delta)$ and T_0^i respectively, to estimate the lapse of time until the next position change of the node requesting the certificate. In other words, the assumption here is that network conditions for nodes that are close to each other should change in a synchronous way. However, due to the asynchronous nature of embedding steps at different nodes, their respective coordinates will not all change at the same time, but we take as "reference" time, the moment when the Surveyor is expected to see a change in its coordinate.

In section 2.3, we showed that most nodes follow a lognormal or Weibull distribution. Depending on which distribution each surveyor node is observing (computing the likelihood at each embedding step), the above formula have a simple form for these two distributions.

For lognormal Inter-Shift distribution we will have:

$$\begin{aligned} S(\Delta) &= 1 - \Phi\left(\frac{\ln \Delta}{\sigma}\right) \\ \mathbb{P}\text{rob}\{T < \tau + \Delta | T > \tau\} &= 1 - \frac{1 - \Phi\left(\frac{\ln(\tau + \Delta - T_0^i)}{\sigma}\right)}{1 - \Phi\left(\frac{\ln(\tau - T_0^i)}{\sigma}\right)} \end{aligned}$$

where the function $\Phi(\cdot)$ is the complementary error function and σ is its shape parameter of the lognormal distribution.

For Weibull distributions we will have:

$$\begin{aligned} S(\Delta) &= 1 - \exp(-\Delta^\gamma) \\ \mathbb{P}\text{rob}\{T < \tau + \Delta | T > \tau\} &= 1 - \frac{1 - \exp(-(\tau + \Delta - T_0^i)^\gamma)}{1 - \exp(-(\tau - T_0^i)^\gamma)} \end{aligned}$$

where γ is the shape parameter of the Weibull distribution.

The validity time of a certificate, Δ , is then computed by setting $\mathbb{P}\text{rob}\{T < \tau + \Delta | T > \tau\} = p_{th}$, where p_{th} is the chosen reliability of the estimate (i.e. the long-term proportions of validity periods that will expire before the corresponding coordinate change).

The certificate then consists in the verified coordinate, the timestamp of the certificate creation, the validity period Δ of the certificate, as well as the identification of the node the certificate is delivered to (i.e. IP address). The certificate is then signed by the the certifying Surveyor using appropriate credentials and encryption keys and issued to the node.

3.2.2 Evaluation

To evaluate the effectiveness of our certificate validity computation, we study the over-estimation resulting from each certificate: for each certificate issued, if the corresponding node moves before its current certificate expires, we record the residual time left on the certificate. This over-estimation is an important security parameter, as it is the lapse of time a malicious node could use a "stale" certificate.

Figure 12 shows the CDF of over-estimation times when the certificates are issued by either the closest Surveyor or a Surveyor chosen at random. We can clearly see that most certificates will not outlive their corresponding coordinates by more than 2 embedding periods in the case where they bare delivered by the closest Surveyor. This is a good result, as many a time, coordinate changes in such timescale will be "localized" in space. The figure also confirms that the accuracy of the validity periods, as thus the security of the system, is improved if coordinates are certified by nearby Surveyors.

Scalability is also an important factor for any certification scheme. Although under-estimation of the validity period of certificates does not pose any security issue for the system, it does tend to increase the load on the system, and on the Surveyors in particular. Obviously, the higher the probability threshold used to compute the certificate validity time, the shorter this time will be. We therefore measure the mean validity period over all certificates for various values of the probability threshold. Figure 13 shows the corresponding average certification rate, which is the inverse of the mean validity period. The average certification rate gives the average number of certification requests that will be issued per node and per time unit (here the embedding period) to the system. This number, multiplied by the number of nodes in the system, and divided by the number of Surveyors and the embedding period gives a certificate request rate per second at each Surveyor.

Figure 13 shows that the average certification rate increases

gently as the probability threshold increases (i.e. as the computation becomes more conservative). This behaviour shows that we can afford a reliability of 95% (and therefore high security) with moderately low overhead.




Figure 12: CDF of Over Estimation Times ($p_{th} = 0.95$), Embedding period = 10mn

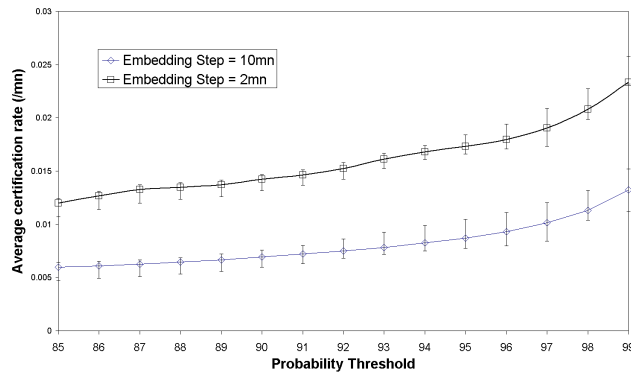


Figure 13: Average Certification Rate

4. DISTANCE ESTIMATION PERFORMANCE EVALUATION

In this section, we study the impact of an attack on the accuracy of distance estimations, with and without our proposed coordinate certification defense mechanism.

The attack consists in malicious nodes choosing, once the Vivaldi system has stabilized, a target they wish to get closer to (each malicious node chooses a target at random amongst the honest nodes), and moving in a straight line towards this target by steps computed from the innovation process of their Kalman filter (see section 3.1.3 and [8] for details). After each displacement, the malicious node seeks to have its newly faked coordinate certified. This strategy is designed to outsmart the tests used for coordinate verification, by only taking (fake) steps that could be attributed to normal system dynamics, as well as caused by normal noise, by the detection test [8]. We choose such subtle attack, as more obvious

displacement caused by fake coordinate would be easier to detect.

Obviously, in the case where coordinate certification is employed, malicious nodes are limited to fake coordinates they can get a certificate for.

We carry out this attack on our PlanetLab experimental setup, with a varying population of malicious nodes.

To assess the impact on distance estimation, we define the following metrics:

- Relative Estimation Error $RER = \frac{|||C'_i - C_j|| - ||C_i - C_j|||}{||C_i - C_j||}$, where C_i is a node's real coordinate in the system without malicious activity, and C'_i is a node's advertised coordinate (and C'_i is either faked, certified or both).
- Security Gain Ratio $SGR = meanRER_{on} / meanRER_{off}$, where $meanRER_{on}$ ($meanRER_{off}$ resp.) is the average RER measured between all pairs of node when the security mechanism is on (off resp.).

Figure 14 shows the SGR observed at the end of the experiment that was allowed to run for a considerable time after convergence of the original (clean) Vivaldi system. The curve labeled "ratio1" depicts the SGR measured in the presence of malicious nodes, while the curve labeled "ratio2" depicts the SGR measured in the clean Vivaldi system without malicious activity.

From this figure, we can conclude that the accuracy of distance estimation in the system with malicious nodes is much improved when coordinate certification is in use than when it is not. This is because the coordinate verification phase of the certification filters out most of the faked displacements. We also see that the curve "ratio2" exhibits a value of 1, indicating that the presence of the certification system does not degrade the performance of the clean system without malicious nodes (in other words, the coordinate certification is very much non intrusive).




Figure 14: Progressive Displacement Attack: Security gain ratio varying the malicious nodes percentage in the overall population.

5. CONCLUSIONS

In this paper we have presented a coordinate certification method to protect the distance estimation phase of Internet Coordinate Systems from nodes lying about their coordinate. This work thus complements our previous work on Coordinate System embedding phase security [8].

The proposed certification method is based on a coordinate verification method, along with a validity period estimation for the corresponding certificate. This method has been shown to be effective, enjoying good verification test performance (high true positive detection rate with low false positive rates), while achieving a very good trade-off between scalability and security. Indeed, the validity periods of certificates are rarely over-estimated, while they still do not trigger too frequent re-certifications.

The reader should note though, that a node knows when its next embedding will occur, and thus when its coordinate is likely to change. A malicious node could exploit this knowledge to seek to obtain a certificate (for its current coordinate) just before performing this embedding, as this could leave such node in possession of a certificate for a soon to be outdated coordinate. To palliate this problem, one could envisage that Surveyors carry out “spot check” on the validity of a node’s certificate: if the certified coordinate fails a new coordinate verification test, the node is “penalized” for using an outdated certificate.

Although this paper focused on Vivaldi for measurements and experimentations, the method proposed for coordinate certification is independent of the embedding protocol used. This because the malicious embedding neighbour detection test that forms the basis of the coordinate verification is itself independent of the specifics of the embedding protocol, and because the validity period computation only depends on observed coordinate inter-shift times. Our proposed method would then be general enough to be applied in the context of coordinates computed by other Internet coordinate systems than Vivaldi.

Finally, regardless of the accuracy of Internet coordinate systems, we believe that securing them is a necessary condition to their deployment. Indeed, without security of both the embedding and distance estimation phases, simple attacks have been shown to easily and seriously disrupt these systems and reduce their utility to next to nothing. This paper thus complements previous work on embedding phase security [8] in this important security mission.

6. REFERENCES

- [1] T. E. Ng, and H. Zhang, *Predicting internet network distance with coordinates-based approaches*, in Proceedings of the IEEE INFOCOM, New York, June 2002.
- [2] M. Pias, J. Crowcroft, S. Wilbur, S. Bhatti, and T. Harris, *Lighthouses for Scalable Distributed Location*, in Proceedings of International Workshop on Peer-to-Peer Systems (IPTPS), Berkeley, February 2003.
- [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), Tokyo, March 2004.
- [4] T. E. Ng and H. Zhang, *A Network Positioning System for the Internet*, in Proceedings of the USENIX annual technical conference, Boston, June 2004.
- [5] F. Dabek, R. Cox, F. Kaashoek and R. Morris, *Vivaldi: A decentralized network coordinate system*. In Proceedings of the ACM SIGCOMM, Portland, Oregon, August 2004.
- [6] Y. Shavitt and T. Tankel, *Big-bang simulation for embedding network distances in euclidean space*, in Proceedings of the IEEE INFOCOM, San Francisco, April 2003.
- [7] R. Zhang, Y. C. Hu, X. Lin, and S. Fahmy, *A Hierarchical Approach to Internet Distance Prediction*, In Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS), Lisboa, July 2006.
- [8] M. A. Kaafar, L. Mathy, C. Barakat, K. Salamatian, T. Turetli and W. Dabbous, *Securing Internet Coordinates Embedding Systems*, in SIGCOMM 2007.
- [9] L. Mathy et al., *Impact of Simple Cheating in Application-Level Multicast*, In Proceedings of Infocom 2003.
- [10] A. Walters, D. Zage and C. Nita-Rotaru, *Mitigating Attacks Against Measurement-Based Adaptation Mechanisms in Unstructured Multicast Overlay Networks*, In Proceedings of ICNP 2006.
- [11] R. Zhang, C. Tang, Y. C. Hu, S. Fahmy and X. Lin, *Impact of the Inaccuracy of Distance Prediction Algorithms on Internet Applications, an Analytical and Comparative Study*, In proceedings of Infocom 2006, Barcelona, April 2006.
- [12] J. Ledlie, P. Gardner, and M. Seltzer, *Network Coordinates in the Wild*, In Proceedings of NSDI, Cambridge, April 2007.
- [13] A. W. Moore et al., *Median Filtering for Removal of Low-Frequency Drift*, Analytic Chemistry, pp. 65:188, 1993.
- [14] Y. Zhang, N. Duffield, V. Paxson, and S. Shenker, *On the Constancy of Internet Path Properties*, in Proceedings of ACM SIGCOMM Internet Measurement Workshop, San Francisco, November 2001.
- [15] G. R. Shorack and J. A. Wellner, *Empirical Processes With Applications to Statistics*, John Wiley&Sons Inc, 1986, p.976.
- [16] H. Zheng, E. K. Lua, M. Pias, and T. Griffin, *Internet Routing Policies and Round-Trip Times*, in Proceedings of the Passive Active Measurement (PAM), Boston, March 2005.

- [17] E. K. Lua, T. griffin, M. Pias, H. Zheng, and J. Crowcroft, *On the accuracy of Embeddings for Internet Coordinate Systems*, in Proceedings of Internet Measurement Conference (IMC), Berkeley, October 2005.
- [18] M. de Berg, et al. *Computational Geometry, Algorithms and Applications*, Springer 2000.