

Securing Internet Coordinate Embedding Systems

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

This paper addresses the issue of the security of Internet Coordinate Systems, by proposing a general method for malicious behavior detection during coordinate computations. We first show that the dynamics of a node, in a coordinate system without abnormal or malicious behavior, can be modeled by a Linear State Space model and tracked by a Kalman filter. Then we show that the obtained model can be generalized in the sense that the parameters of a filter calibrated at a node can be used effectively to model and predict the dynamic behavior at another node, as long as the two nodes are not too far apart in the network. This leads to the proposal of a Surveyor infrastructure: Surveyor nodes are trusted, honest nodes that use each other exclusively to position themselves in the coordinate space, and are therefore immune to malicious behavior in the system. During their own coordinate embedding, other nodes can then use the filter parameters of a nearby Surveyor as a representation of normal, clean system behavior to detect and filter out abnormal or malicious activity. A combination of simulations and Planet-Lab experiments are used to demonstrate the validity, generality, and effectiveness of the proposed approach for two representative coordinate embedding systems, namely Vivaldi and NPS.

Categories and Subject Descriptors

C.2 [Computer-Communications Networks]: Security and protection; C.2.1 [Network Architecture and Design]: Network topology

General Terms

Security, Performance, Experimentation

Keywords

Network positioning systems, Internet Coordinates-embedding systems, security, malicious behavior detection, Kalman filter

^{*}L. Mathy was also partly supported by LAAS-CNRS, France and the University of Liège, Belgium.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGCOMM'07, August 27–31, 2007, Kyoto, Japan.
Copyright 2007 ACM 978-1-59593-713-1/07/0008 ...\$5.00.

1. INTRODUCTION

Internet coordinate systems (e.g., [1, 2, 3, 4, 5, 6], etc.) embed latency measurements amongst samples of a node population into a geometric space and associate a network coordinate vector (or coordinate in short) in this geometric space to each node, with a view to enable accurate and cheap distance (i.e. latency) predictions amongst any pair of nodes in the population. Extensive measurements and analysis from a live, large-scale deployment have shown network coordinate systems to be fit for purpose [7], making them a valuable tool to support distributed applications, systems and overlays (e.g., [8, 9, 10]) that rely on, and benefit from, the notion of network topology-awareness. However, it should also be noted that coordinate-based positioning systems only achieve desirable accuracy, robustness, stability and scalability properties at the expense of rather slow convergence times – in other words, a new node joining the coordinate system may not reach an accurate value for its own coordinate before several tens of seconds or even several minutes. Such convergence properties would argue in favor of a deployment of Internet coordinate systems as an always-on, large-scale service.

Unfortunately, such coordinate services have been shown to be vulnerable to malicious attacks, providing a potentially attractive fertile ground for the disruption or collapse of the many applications and overlays that would use these services [11]. There are actually two obvious ways to disrupt the operation of a coordinate based system. First when requested to give its coordinate for a distance estimation at the application-level, a malicious node could simply and blatantly lie. Second, a malicious node, or even a colluding group, may aim at disrupting the embedding process itself. This latter strategy is very insidious and effective as it can result in important distortions of the coordinate space which then spoils the coordinate computations of many nodes (malicious and honest alike) [11]. This paper focuses on developing and studying generic methods to secure the coordinate embedding process (the problem of nodes lying about their coordinate during distance estimation at the application-level is not addressed).

More precisely, the embedding process, regardless of the actual coordinate-based positioning system, works on the premise that nodes adjust their coordinate based on some comparison between measured and estimated distances to some other nodes. Malicious nodes can interfere with this embedding process by, amongst other things, lying about their real coordinate and/or tampering with measurement probes, to create a discrepancy between measured and estimated latencies, so that unsuspecting nodes would wrongly adjust their own coordinate in a bid to reduce the difference [12]. Because the load on the network naturally varies in time, so does latency between a pair of nodes, and as a result, the embedding process must

be run periodically by all nodes to track changes in network conditions. This “continuous” adjustment of nodes’ coordinates can not only result in a drift of the coordinate space [7] but also gives plenty of scope and opportunities for malicious activity. We therefore seek to equip (honest) nodes with a means to detect, with low overhead, malicious activities they may encounter during embedding.

Noting that, in the absence of malicious nodes, a node’s coordinate depends on the combination of network conditions and the specificities of the embedding process itself (e.g. which coordinate protocol is in use, the chosen dimensionality of the geometric space, etc), we therefore introduce the concept of *Surveyor nodes* (or *Surveyors* in short). Surveyors form a group of trusted (honest) nodes, scattered across the network, which use each other *exclusively* to position themselves in the coordinate space. Of course, Surveyors do assist other nodes in their positioning (as prescribed by the embedding protocol), but we stress that Surveyors never rely on non-Surveyor nodes to compute their own coordinate. This strategy thus allows Surveyors to experience and learn the natural evolution of the coordinate space, as observed by the evolution of their own coordinate, in the absence of malicious activities. In essence, Surveyor nodes are thus vantage points guaranteed to be immune from malicious activities. The idea is that Surveyors can then share a “representation” of normal behavior in the system with other nodes to enable them to detect and filter out abnormal behavior.

We postulate and verify that, in the absence of malicious activity, a node’s coordinate can be viewed as a stochastic process with linear dependencies whose evolution can be tracked by a Kalman filter [13, 14]. Each Surveyor then computes and calibrates the parameters of a linear state space model and shares the parameters of this model with other nodes. These nodes can then use these parameters, to run locally and in a “stand-alone” fashion a Kalman filter tracking the coordinate adjustments. These nodes can then use the Kalman filter output (the innovation process), to compare their observed coordinate adjustments with the one predicted by the Kalman filter, and flag as “suspicious” embedding steps where the difference would be too high.

In section 2, we present a general model of coordinate embedding, in the absence of malicious nodes, that naturally leads to the Kalman filter framework. In section 3, we validate the model, with both simulations and PlanetLab experiments, in the case of both Vivaldi [5] and NPS [4]. This section also studies the viability of the idea of using Surveyor nodes in secure coordinate embedding. We then describe and evaluate, in sections 4 and 5, how Surveyors can effectively be used for malicious node detection in the specific embedding process of Vivaldi and NPS. We conclude with perspectives on our contributions in section 6.

2. COORDINATE EMBEDDING MODEL

The goal of embedding systems, regardless of the embedding method and geometric space used, is to assign a coordinate to every node in the system so that, at any time, the distance between any two points in the geometric space should provide a good estimate of the network distance, measured as an RTT (Round Trip Time), between the corresponding nodes. Obviously, 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) as well as the state of the operating system in nodes (e.g. scheduling state generating measurement noise, 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 [15], 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 coordinate on a periodic basis. Consequently, the coordinate of a node can be viewed as a discrete stochastic process, and we will use X_i^n to represent the coordinate of node i at “discrete time” n .

Without loss of generality, consider that a node (called the embedding node) computes its coordinate through a series of embedding steps, where each embedding step represents a coordinate adjustment based on a one-to-one interaction with another node, called a peer node (e.g. peer nodes are called neighbors in Vivaldi, and landmarks or reference points in NPS). Note that when the embedding protocol requires that a node uses several peer nodes simultaneously for repositioning, for the purpose of our modelization, we simply consider that each peer node corresponds to a distinct embedding step, each taking place at “successive” discrete times.

At every embedding step, the “fitness” (or “correctness”) of the embedding node coordinate is assessed by computing the deviation between the measured RTT towards the corresponding peer node and the one estimated in the coordinate system. More precisely, suppose that at its n^{th} embedding step, embedding node i has current coordinate X_i^n and uses peer node j with current coordinate X_j^n . Suppose that the RTT between these nodes, measured during this embedding step, is RTT_{ij}^n . The fitness of the embedding node coordinate can then be computed as the *measured relative error* $D_n = \frac{||X_i^n - X_j^n|| - RTT_{ij}^n}{RTT_{ij}^n}$. The goal of *any* embedding system, regardless of the embedding method proposed and/or the geometric space structure, is to minimize a “cost” indicator (e.g. mean square error) that captures the measured relative error that could be observed between any node and any other node in the system, at any time.

As the measured relative errors are fundamental performance indicators to all embedding systems, it seems natural to develop a model that captures their dynamic characteristics, although we note that relative errors often have complex behavior (and may thus not be a natural choice from a modeling perspective).

Measured relative errors are subject to fluctuations of the RTT for the reasons mentioned above, namely transient network congestion and operating system scheduling issues. To isolate the impact of these RTT fluctuations on anomaly detection, we introduce Δ_n , the *nominal relative error* that our node under consideration would have obtained at its n^{th} embedding step if the RTTs in the network had not fluctuated. An anomaly becomes simply a large deviation of measured relative error D_n from its nominal value defined by Δ_n .

Because many sources contribute to the deviation of D_n from its nominal value (RTT measurement error, RTT fluctuations, errors in node coordinates), it is reasonable to suppose that they relate to each other as follows,

$$D_n = \Delta_n + U_n \quad (1)$$

where U_n is a Gaussian random variable with mean zero and variance v_U .

We now focus on the dynamics of the system in its nominal regime where RTTs do not fluctuate. In the absence of complete and accurate knowledge of the system, nodes keep on adapting the nominal relative error on a pairwise basis with their peer nodes, aiming to optimize the cost indicator. This adaptation is subject to an error caused by the other nodes in the system adapting their coordinate (and corresponding relative error) in a completely distributed way. We thus define the *system error* W_n which represents the impact of other nodes on the positioning of a node at embedding step

n . Since the system error at a node results from many contributing sources, it is also reasonable to assume that it is a white gaussian process (with mean \bar{w} and variance v_W)¹.

Because of the nature of large-scale embedding processes, the nominal relative error Δ_n can be deemed to follow a stochastic process that converges to some stationary regime characterized by a positive average. As a first approximation, the process Δ_n could be modeled as a first order Auto Regressive (AR) model:

$$\Delta_{n+1} = \beta\Delta_n + W_n. \quad (2)$$

where β is a constant factor strictly less than one otherwise the relative error does not converge to a stationary regime independently of the initial condition. This equation captures the dynamic evolution of the nominal relative error of a node through successive embedding steps.

Equations 2 and 1 define a linear state space model for the relative error of a node. Our goal is to devise a way to obtain relative error predictions from this model. Because of the linear properties of the model, a Kalman filter can be used to track the evolution of the nominal relative error and obtain a *predicted relative error* $\hat{\Delta}_{n|n-1}$ (see section 2.1).

The idea behind this strategy is that if the stochastic space model, and especially its associated Kalman filter, are calibrated within a clean embedding system, then a simple hypothesis test can be used to assess whether the deviation between the measured relative error and the predicted relative error, observed at a given embedding step, is normal or is the fruit of anomalous or malicious activity from the peer node. From this perspective, even if the state space model considered is crude, its quality should be evaluated based on the final outcome in terms of probability of detection and false positive rate. We will see in the evaluation section (section 5) that this model achieves very good performance.

2.1 Kalman Filter Equations

The Kalman filter is used here to estimate Δ_n given the set of previously measured relative errors $\mathcal{D}_0^n = \{D_0, \dots, D_n\}$. Under the hypothesis of a gaussian noise process in the underlying state space model, the Kalman filter gives the Least Mean Squared estimates of Δ_n , $\hat{\Delta}_n$. Moreover, it gives the quality of these estimates through an evaluation of the mean squared error i.e., $E[(\hat{\Delta}_n - \Delta_n)^2]$. This last value could be used to detect anomalies through large deviations of the measured relative error from its mean.

We will assume here that all the parameters of the space model given in Eq. (1) and Eq. (2) are known and given. In the next section we will describe how to derive these parameters.

Let us denote by $\hat{\Delta}_{i|i-1}$ the estimation of Δ_i knowing the observations of network delay up to time $i-1$, and $\hat{\Delta}_{i|i}$ the estimate after the measurement D_i is done. Similarly, let $P_{i|i-1}$ be the estimated *a posteriori* error variance at time i knowing the observations up to time $i-1$, and let $P_{i|i}$ be the estimation of the *a posteriori* error variance after D_i is known. The Kalman Filter is composed of two steps that are iterated. The first step is called the prediction step and the second one the update step.

In the prediction step, the value of $\hat{\Delta}_{i|i-1}$ is calculated based on $\hat{\Delta}_{i-1|i-1}$ as :

$$\hat{\Delta}_{i|i-1} = \beta\hat{\Delta}_{i-1|i-1} + \bar{w}.$$

The *a posteriori* error variance of this estimate is :

$$P_{i|i-1} = \beta^2 P_{i-1|i-1} + v_W.$$

¹The value \bar{w} accounts for the drift that has been observed in positioning systems [7].

In the update step, $\hat{\Delta}_{i|i-1}$ is updated to integrate the observed measurement D_i :

$$\hat{\Delta}_{i|i} = \hat{\Delta}_{i|i-1} + K_i(D_i - \hat{\Delta}_{i|i-1})$$

where K_i denotes the updated gain and is obtained as :

$$K_i = \frac{P_{i|i-1}}{P_{i|i-1} + v_U}.$$

The *a posteriori* error variance of this estimate is :

$$P_{i|i} = \frac{v_U}{P_{i|i-1} + v_U} P_{i|i-1}.$$

The value $\eta_i = D_i - \hat{\Delta}_{i|i-1}$ is called the innovation process and is the main process to observe for anomalous behavior detection (see section 4.1). The innovation process is a white (meaning that it is an independent process) gaussian process with a mean 0 and a variance equal to $v_{\eta,i} = v_U + P_{i|i-1}$. Abnormality simply amounts to a significant deviation from the nominal values of the innovation process characterized by the Kalman filter.

To run the Kalman estimation, we need as initial values the system state value w_0 and the *a priori* state variance $P_{0|0} = p_0$. These two values are estimated during the parameters calibration step.

2.2 Calibration of the Kalman filter

Before running the estimation using the Kalman filter, the values of the filter parameters $\theta = (\beta, v_W, v_U, w_0, p_0)$ have to be computed. For this purpose we need to calibrate these parameters over coordinate measurements collected during a stationary and cheater-free period. The calibration can be done using a maximum likelihood criteria (choosing parameter values such that the likelihood of observing the measurements is maximized) by applying the Expectation Maximization (EM) method. We follow the approach presented in [16] for the EM derivation.

Calibration by EM method.

Let's assume that \mathcal{D}_0^N is the set of all measured prediction errors, $\mathcal{D}_0^N = \{D_0, \dots, D_N\}$ and let $\Delta_0^N = \{\Delta_0, \dots, \Delta_N\}$ be the set of nominal relative errors.

As all the noise processes are assumed to be gaussian, \mathcal{D}_0^N and Δ_0^N will jointly follow a gaussian distribution. The log-likelihood of \mathcal{D}_0^N and Δ_0^N can therefore be written as follows:

$$\begin{aligned} \log \text{Prob}\{\mathcal{D}_0^N, \Delta_0^N\} &= - \sum_{i=0}^N \frac{(D_i - \Delta_i)^2}{2v_U} \\ &- \frac{N+1}{2} \log v_U - \sum_{i=1}^N \frac{(\Delta_i - \beta\Delta_{i-1} - \bar{w})^2}{2v_W} \\ &- \frac{N}{2} \log v_W - \frac{(\Delta_0 - w_0)^2}{2p_0} - \frac{1}{2} \log p_0 \\ &- (N+1) \log 2\pi. \end{aligned}$$

In a Maximum likelihood setting, we wish to find the values for the parameters that will maximize the above log-likelihood assuming that the sequence \mathcal{D}_0^N has been observed. However as the sequence of system state Δ_0^N has not been observed, this maximization is not tractable directly and we have to apply the Expectation Maximization method. This method transforms the maximization of the above likelihood function with unobserved system state sequence Δ_0^N to an iteration of successive steps where the system state sequence is assumed to be known and the parameters can be obtained through maximization of the likelihood function.

Each iteration of the EM method consists therefore of two steps. In the first step, we compute the expectation (over all possible values of the sequence of states Δ_0^N) of the log-likelihood, given the observed values of D_n and assume that the parameter values are equal to $\theta^{(k)}$. In a second step, the parameters $\theta^{(k+1)}$ are chosen so as to maximize the previously obtained likelihood expectation. Next we explain these two operations with some further details.

Let the superscript (k) indicate the value of any parameter at the k^{th} step of the EM algorithm. As explained before, in the EM method, we need to estimate the value of the unobserved system states to be able to calculate the overall likelihood to maximize. The variables $\hat{\delta}_i^{(k)}$ are in fact those estimates at iteration k and $\hat{\pi}_i^{(k)}$ and $\hat{\pi}_{i,i-1}^{(k)}$ are the estimation error variances of this sequence of states:

$$\hat{\delta}_i^{(k)} = E[\Delta_i | \mathcal{D}_0^N, \theta^{(k)}], \quad \hat{\pi}_i^{(k)} = E[\Delta_i^2 | \mathcal{D}_0^N, \theta^{(k)}],$$

$$\hat{\pi}_{i,i-1}^{(k)} = E[\Delta_i \Delta_{i-1} | \mathcal{D}_0^N, \theta^{(k)}].$$

Expectation step.

The expected value of log-likelihood knowing the set of measured values \mathcal{D}_0^N and the parameter $\theta^{(k)}$ is given by :

$$\begin{aligned} \bar{L}(\theta, \theta^{(k)}) &= E[\log \text{Prob}\{\mathcal{D}_0^N, \Delta_0^N\} | \mathcal{D}_0^N, \theta^{(k)}] \\ &= - \sum_{i=0}^N \frac{D_i^2 - 2D_i \hat{\delta}_i^{(k)} + \hat{\pi}_i^{(k)}}{2v_U} \\ &\quad - \frac{N+1}{2} \log v_U - \sum_{i=1}^N \frac{\hat{\pi}_i^{(k)} + \beta^2 \hat{\pi}_{i-1}^{(k)} + \bar{w}^2}{2v_W} \\ &\quad - \sum_{i=1}^N \frac{\beta \hat{\pi}_{i,i-1}^{(k)} + \hat{\delta}_i^{(k)} \bar{w} + \beta \hat{\delta}_{i-1}^{(k)} \bar{w}}{v_W} \\ &\quad - \frac{N}{2} \log v_W - \frac{p_0 - 2\hat{\delta}_0^{(k)} w_0 + w_0^2}{2p_0} \\ &\quad - \frac{1}{2} \log p_0 - (N+1) \log 2\pi. \end{aligned}$$

By replacing θ by its value at the k^{th} step of the EM algorithm, we obtain $\hat{\delta}_i^{(k)}$, $\hat{\pi}_i^{(k)}$ and $\hat{\pi}_{i,i-1}^{(k)}$, which gives the expected log-likelihood at the $(k+1)^{\text{th}}$ step. Interested readers can find details of the computations of these values in [17].

Maximization step.

In this step, the parameter vector at step $(k+1)$ is chosen to maximize the expected log-likelihood. This is done by solving the equation

$$\frac{\partial \bar{L}(\theta, \theta^{(k)})}{\partial \theta} = 0.$$

This results in the following set of equations:

$$\begin{cases} w_0^{(k+1)} = \hat{\delta}_0^{(k)} \\ p_0^{(k+1)} = p_0^{(k)} - (\hat{\delta}_0^{(k)})^2 \\ v_U^{(k+1)} = \frac{1}{N+1} \sum_{i=0}^N D_i^2 - 2D_i \hat{\delta}_i^{(k)} + \hat{\pi}_i^{(k)} \\ \bar{w}^{(k+1)} = \sum_{i=1}^N \hat{\delta}_i^{(k)} + \beta^{(k+1)} \sum_{i=1}^N \hat{\delta}_{i-1}^{(k)} \\ \beta^{(k+1)} = \frac{\sum_{i=1}^N \hat{\pi}_{i,i-1}^{(k)} + \bar{w}^{(k+1)} \sum_{i=1}^N \hat{\delta}_{i-1}^{(k)}}{\sum_{i=1}^N \hat{\pi}_{i-1}^{(k)}} \\ v_W^{(k+1)} = \frac{1}{N} \sum_{i=1}^N \hat{\pi}_i^{(k)} + (\beta^{(k+1)})^2 \hat{\pi}_{i-1}^{(k)} + (\bar{w}^{(k+1)})^2 \\ - 2\beta^{(k+1)} \hat{\pi}_{i,i-1}^{(k)} - 2\hat{\delta}_i^{(k)} \bar{w}^{(k+1)} - 2\beta^{(k+1)} \hat{\delta}_{i-1}^{(k)} \bar{w}^{(k+1)}. \end{cases}$$

By solving this set of equations we can obtain the vector $\theta^{(k+1)}$, then we iterate with the expectation calculation as described above.

We note that the complexity of the approach lies in the linear state space modeling phase by EM algorithm that incurs a number of iterations over N dimensional vectors, which is well within the capability of modern computers. We will see later that this phase has to be run on a subset of nodes (the Surveyors). On the other hand, predicting relative errors using the Kalman filter (section 2.1), which occurs on every node, only implies a few simple scalar operations and is negligible in terms of required computing power.

Finally, because we expect each of the innovation observation η_n to be inside a confidence interval of $\pm 2\sqrt{v_{\eta,n}}$ (where $v_{\eta,n}$ is the variance of the innovation process at time n) with a probability higher than 95%, when a Kalman filter yields 10 consecutive innovation observations outside such confidence interval, the filter is re-calibrated by re-applying the calibration procedure described in this section. Re-calibration is likely to occur following a significant change in the corresponding node's coordinates, caused by changes in network conditions.

3. VALIDATION

To validate our model, we conducted simulations and PlanetLab experiments for both Vivaldi [5] and NPS [4]. Vivaldi is a prominent representative of purely peer-to-peer-based (i.e. without infrastructure support) positioning systems, where the system tries to minimize the potential energy of virtual springs between nodes, while NPS is typical of infrastructure-based systems, where a hierarchy of landmarks and reference points governs the positioning of nodes.

As the goal of this section is to assess the fitness of the proposed model to represent the normal behavior of the embedding processes, all results presented were acquired in a clean environment with no malicious node. While the goal of the simulation studies is to assess our results for large scale coordinate-based systems, the PlanetLab experiments aim to show their applicability in real-world conditions.

The PlanetLab experiments were conducted over a set of 280 PlanetLab nodes spread world-wide. In this paper, we discuss a representative set of experimental results conducted over several days in December 2006.

The simulations were driven by a matrix of inter-host Internet RTTs (the ‘‘King’’ dataset) to model latencies based on real world measurements. This dataset contains the pair-wise RTTs between 1740 Internet DNS servers collected using the King method [18] and was used to generate a topology with 1740 overlay nodes.

In the case of Vivaldi, each node had 64 neighbors (i.e. was attached to 64 springs), 32 of which being chosen to be closer than 50 ms. The constant fraction C_e for the adaptive timestep (see [5] for details) is set to 0.25. A 2-dimensional coordinate space augmented with a height vector was used.

For NPS, we considered an 8-dimensional Euclidean space for the embedding. We used an NPS positioning hierarchy with 4 layers. The top layer had a set of 20 well separated permanent landmarks. Each subsequent layer then had 20% of nodes randomly chosen as reference points. The security mechanism already proposed in NPS, shown to be too primitive in [11], was turned on and its sensitivity was set to 4 (see [4] for details).

When needed, Surveyor nodes were chosen at random².

²Note that in NPS, all permanent landmarks also act as Surveyors.

3.1 Assumption Validation

In section 2, the assumption that the system error W_n follows a gaussian distribution was made. This is fundamental to the applicability of the Kalman filter framework. Every node calibrated its own Kalman filter based on the observation of its own embedding, and we checked this assumption by applying the Lillie test [19], a robust version of the well known kolmogoroff-Smirnov goodness-of-fit test, to whitened filter inputs. We observed that the Lillie test leads to only 14 gaussian fitting rejections in simulations (over 1720 samples) and 5 rejections in PlanetLab (over 260 samples). This test allows us to conclude that the hypothesis we took for the Kalman model is valid. In addition, we plot in figure 1 the Quantile-

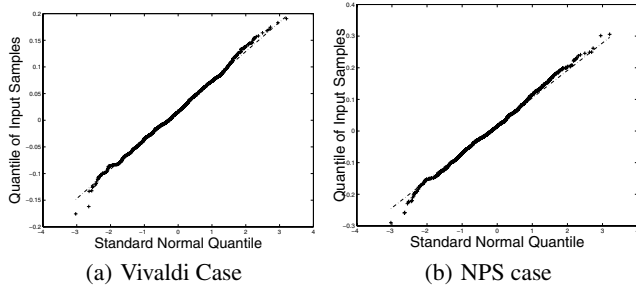


Figure 1: Quantile-Quantile plot of 2 innovation processes.

Quantile (QQ) plots of 2 innovation processes (for both Vivaldi and NPS) taken on PlanetLab nodes running their own Kalman filter. These plots, typical of observations on all nodes, show that each of these distributions indeed closely follows a Gaussian distribution.

3.2 Effective Behavior Representation

From section 2, it is clear that the Kalman filter model attempts to represent the behavior of the embedding process by capturing the dynamics of the system through its convergence behavior (by tracking of relative errors over time). In this section, we therefore assess the representational power of this approach by having each node calibrate its own Kalman filter from the measurements it observed during the embedding of its own coordinates, in a cheat-free regime. Then, once the model has converged at every node (i.e. the EM method has converged and the variations of all the θ components become smaller than 0.02), a new embedding process is started (i.e. the nodes forget their coordinates and rejoin the system). During this second embedding process, the prediction error, that is the absolute value between the error predicted by the node’s Kalman filter and the measured actual error, is measured.

Figure 2 shows a typical evolution of actual (measured) relative errors and predicted errors for a node on PlanetLab (for Vivaldi, but similar behavior was observed for NPS). The two curves of the top graph of the figure are so similar that they are almost indistinguishable. The bottom graph of the figure represents the prediction error which is the difference between these two curves (note the different scale used for this graph). We see that the prediction errors are small which shows that a node’s calibrated Kalman filter can capture effectively the node’s behavior “in the wild”.

Figure 3 shows the cumulative distribution function (CDF) of all the prediction errors observed across all nodes in the system. This confirms that the vast majority of predictions are indeed excellent. This demonstrates the power and generality of the model in capturing the dynamics of the system and its adaptability to current system conditions.

However, there are a few “outlier” predictions with large errors.




Figure 2: Prediction errors (PlanetLab node).




Figure 3: CDF of prediction errors.

To better understand their nature, we show in table 1 the repartition of prediction errors in error intervals. The table shows the number of nodes with prediction errors in this interval, the number of occurrences of the smallest prediction error observed in the interval, and the number of occurrences of the largest prediction errors observed in the interval (e.g. Number of node(s)/number of observed min error/number of observed max error). We see that only a few nodes contribute (sometimes very many) large prediction errors. Looking further, we identified 3 nodes, all located in India, who contributed consistently to the “tail” of the CDF in figure 3. It is interesting to note that these nodes exhibited large (> 0.75) average measured relative errors during embedding, and were clearly having trouble with the embedding process itself, due to adverse network conditions.

Table 1: Prediction Error Histogram

Error Interval	(Vivaldi)	(NPS)
0.0-0.05	257/830/922	232/854/943
0.05-0.1	32/201/995	44/180/941
0.1-0.15	5/3/992	18/5/229
0.15-0.2	1/997/997	2/12/884
0.4-0.45	4/3/56	3/5/12
0.45-0.5	1/12/12	2/1/17
0.5-0.55	1/985/985	2/32/40
0.6-0.65	-	1/851/851

3.3 Representativeness of Surveyors

If a subset of nodes in the system (called Surveyor nodes) are trusted and use each other exclusively to compute their own coordinates, they will be immune to the effects of malicious behavior during embedding. The premise of our work is that the “clean” (normal) system behavior thus learnt can then be shared with other nodes and used by these nodes to detect malicious behavior they may be subjected to by other nodes in the system. This obviously assumes that the behavior of the system as observed by Surveyors can approximate or represent well enough the normal behavior of the system as observed by other nodes in the absence of malicious behavior. In the following validation of this assumption, Surveyors are chosen at random in the node population.

Note that a random choice will give an upper bound on the number of Surveyors needed. Indeed, intuitively, Surveyors should be roughly uniformly distributed in the system to be representative of most other nodes. However, choosing nodes at random in the system does not yield a uniform distribution of Surveyors (i.e. “Surveyor clusters” appear) and therefore not every new Surveyor increases representativeness. Consequently, more Surveyors must be chosen in order to achieve a good coverage of the system, than if they were placed more strategically. Nevertheless, the random choice method does provide general results, without the need to address the question of optimal Surveyor deployment strategy (which we leave as future work).

One of the first questions to answer is how many Surveyors are needed to be representative of the rest of the population. Noting that our model is based on measured relative errors and that each node in the system observes a series (i.e. distribution) of such errors, we characterize the system-wide relative error behavior as the CDF of the 95th percentiles of the relative errors observed at each (normal) node (i.e. the distribution is made up of the 95th percentile value observed at every node). We then compare this CDF with those of the 95th percentiles of the relative errors observed across a varying population of Surveyors. The choice of the 95th percentile is so that outliers, as observed in section 3.2, do not skew the results, while preserving a high degree of generality. Figure 4,




Figure 4: Impact of Surveyor population size on representativeness.

obtained by simulations of Vivaldi, indicates that a population of Surveyors of about 8% of the overall population is closely representative of this overall population (because the CDF for these populations in the figure are similar).

To generalize this result, we then repeated the experiment, using both simulations and PlanetLab measurements, on a Vivaldi system with 8% of Surveyors. We again chose to represent the system by the distribution of the 95th percentile of the measured relative errors (figure 5).

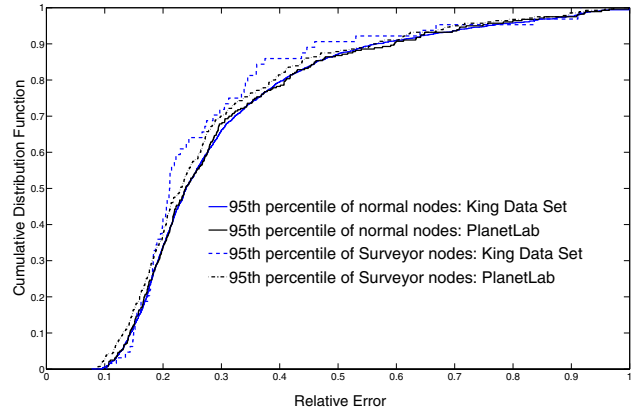


Figure 5: Representativeness with 8% Surveyor nodes.

These experiments confirm that less than 10% of randomly chosen Surveyor nodes is enough to gain a good representation of the system behavior. Similar results were observed for NPS. We note that 8% to 10% of the overall node population is a very stringent requirement for most practical purposes and can represent a huge number of nodes. However, as already pointed out above, random Surveyor deployment is not optimal and this value is an upper bound on the number of Surveyors needed. To gain further insight into how conservative this upper bound may be, we tried a simple k-means clustering algorithm for Surveyor deployment. Figure 4 shows that when taking cluster heads as Surveyors, good representativeness can be achieved with roughly 1% of Surveyors. Although this does not give much indication as to what the lower bound on the number of Surveyors needed is in the case of optimal Surveyor deployment, it nevertheless shows that simple deployment methods can reduce requirements considerably and that the upper bound yielded by random deployment is indeed very conservative.

Having shown that a population of Surveyors can represent the overall system, the next question is how well the behavior of the system as captured by the Kalman filter calibrated by a Surveyor, can represent the behavior of a single (normal) node. To answer this question, we carried out an experiment where a population of nodes took part in a Vivaldi embedding on PlanetLab. Each node used the Kalman filter of every Surveyor and generated multiple prediction errors (one per Surveyor) at every embedding step.

Figure 6 shows the maximum prediction error yielded by each Surveyor, for each normal node in the system, observed during this experiment. What we observe is that although each normal node can find at least one Surveyor node whose Kalman filter yields very low prediction errors, not every Surveyor is a good representative for any given normal node. The Surveyor chosen as a representative by a normal node is therefore important to achieve good prediction performance (and thus good malicious behavior detection).

Figure 7 plots the prediction accuracy (measured as an average prediction error) against the distance (measured as an RTT) between a node and the corresponding Surveyor, as observed during the PlanetLab experiment. It is clear that better locality between a node and its Surveyor yields more accurate predictions. This property seems intuitive, as a Surveyor closer in terms of RTT will also be closer in the geometric space, and will thus be more likely to experience dynamics of the coordinate system similar to that of the local area where the node resides. This is confirmed in figure 8 which shows the maximum prediction error, observed for Vivaldi




Figure 6: Maximum prediction errors with Surveyor filter parameters (PlanetLab).

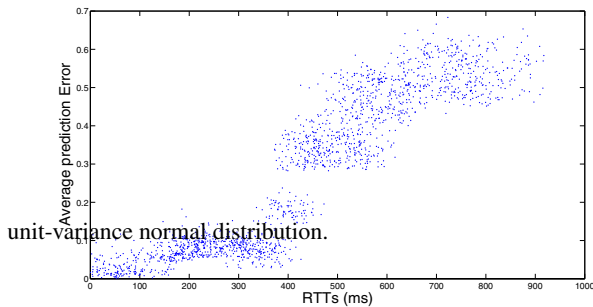


Figure 7: Correlation between 'Node-Surveyor' RTTs and estimation accuracy (PlanetLab).

on PlanetLab, when nodes use the closest Surveyor as their representative. Similar results were observed for PlanetLab experiments with NPS.

Finally, it is again important to note that all the results in this section were obtained with randomly chosen Surveyors. Strategically placing Surveyors to ensure a better coverage of the network and coordinate space, would simply improve the prediction accuracy, while reducing the number of Surveyors required.

4. MALICIOUS BEHAVIOR DETECTION

The previous section has shown that normal node behavior can be modeled by a Kalman filter. More importantly, it has also been shown that this technique is powerful and robust enough that the normal behavior model captured on one node is readily and effectively applicable on other nearby nodes. This property leads to the idea of Surveyors.

Surveyors are a set of nodes in the coordinate space that exclusively use each other to compute their own coordinates. In other words, in Vivaldi, Surveyors only use other Surveyors as neighbors, while in NPS, they only use other Surveyors as reference points (note that in NPS, all landmarks also act as Surveyors, although not all reference points will be Surveyors). Of course, Surveyors can, and will be chosen as neighbors or reference points by other (non-Surveyor) nodes in the system, but the point is that a Surveyor




Figure 8: Maximum prediction errors with closest Surveyor.

adjusts its coordinate solely in response to embedding steps (i.e. measurements) with other Surveyors. If Surveyors run a clean version of the coordinate embedding software and they are carefully kept clean of malicious software, such as viruses or worms, that could implement malicious modifications to the embedding, then they can be considered as clean, honest nodes. Because Surveyors only interact with each other during their own embedding, they are therefore immune to malicious or anomalous behavior in the system, and they therefore observe the behavior of the system in clean, normal conditions. The idea is then to use the thus obtained normal behavior model as a basis for anomalous behavior detection at other nodes of the system. To do so, nodes use the parameters of the Kalman filter calibrated at a nearby Surveyor.

It is important to note that the proposed method is entirely distributed as each node has its own filter. Indeed, Surveyors calibrate and recalibrate their own filter as needed, depending on varying network conditions, and share the resulting filter parameters with other nodes, but they take no further active part in anomalous behavior detection at other nodes. When a node's filter needs re-calibrating (e.g. because it starts giving too many detection alarms), the node simply obtains fresh filter parameters from a Surveyor.

4.1 Anomalous Behavior Detection Method

At each embedding step, a node computes a measured relative error D_n towards a peer node. Recall from section 2 that the Kalman filter at the node can provide $\hat{\Delta}_{n|n-1}$, the predicted relative error from the previously measured relative errors. The innovation process of the Kalman filter yields the deviation between the measured and predicted relative errors, $\eta_n = (D_n - \hat{\Delta}_{n|n-1})$, which, in a system without malicious node, follows a zero-mean gaussian distribution with variance $v_{\eta,n} = v_U + P_{n|n-1}$ (also yielded by the filter).

This allows us to detect malicious behavior as a simple hypothesis test. Let H_0 be the hypothesis that the peer node has a normal behavior (i.e. it is honest). The hypothesis testing simply consists of assessing whether the deviation between the measured and predicted relative errors is normal enough under expected system behavior. Given a "significance level" α , which determines the "aggressivity" or "strictness" of the test, the problem is to find the threshold value t_n such that

$$P(|D_n - \hat{\Delta}_{n|n-1}| \geq t_n | H_0) = \alpha. \quad (3)$$

But since, under hypothesis H_0 , $(D_n - \hat{\Delta}_{n|n-1})$ follows a zero-mean normal distribution with variance $v_{\eta,n}$, we also have that

$$P(|D_n - \hat{\Delta}_{n|n-1}| \geq t_n | H_0) = 2Q(t_n/\sqrt{v_{\eta,n}}), \quad (4)$$

where $Q(x) = 1 - \Phi(x)$, with $\Phi(x)$ being the CDF of a zero-mean, unit-variance normal distribution.

From equations 3 and 4, we therefore have

$$t_n = \sqrt{v_{\eta,n}} Q^{-1}(\alpha/2). \quad (5)$$

If the observed deviation exceeds the threshold given by equation 5, then the hypothesis is rejected, the peer node is flagged as suspicious, the embedding step is aborted and the measured relative error D_n is discarded (i.e. it is not used to update the state of the filter).

Note that a suspicious node, as detected by this test, is not necessarily associated with malicious intent, but could be caused by changing network conditions. Honest nodes classified as suspicious represent false positives and have little impact on the system as long as their occurrence is low. The trade-off between aggressivity and strictness of the test is represented by the so called ROC (Receiver Operation Characteristic) curves [20]. These curves plot the true positive rate versus the false positive rate, *i.e.* the probability of correctly detecting a malicious node versus the probability of labelling an honest node as malicious. In practice trying to increase the true positive rate (the probability of malicious node detection) comes at the cost of increasing the false positive rate.

4.2 Generic Detection protocol

In general, on identifying a peer node as suspicious, a node will replace it, that is choose a new neighbor in Vivaldi or a new reference point in NPS.

The only exception to this rule is when the node was embedding against the peer node for the very first time. In this particular case, the node uses its local error e_l ³ as an indicator of the confidence it has in its own coordinates, to carry out a second hypothesis test identical to that presented in the previous section, but this time with a confidence level of $e_l\alpha$. If the test is accepted, then the peer node gets a reprieve and is not replaced, so that a second embedding against this peer node will be attempted at a later time.

The main idea behind this potential reprieve for first-time peer nodes is that a node whose coordinate has already converged towards its true value can afford a few aborted embedding steps with very little impact on the accuracy of its coordinate. On the other hand, a new peer node which is in the process of joining the network may trigger the abortion of an embedding step, simply because its coordinate has not converged yet (as opposed to because it displays a malicious behavior). In this case, the reprieve simply gives time to the new (joining) peer node to converge before being identified as malicious. Of course an embedding node which is not confident in its coordinate must strive to reduce the number of aborted embedding steps so as not to compromise its convergence in the system, and will therefore grant fewer reprieves (because its e_l is higher) than a node that has already converged.

Finally, we use a simple mechanism for the selection of the Surveyor from which a node obtains its calibrated Kalman filter. All Surveyor nodes register with an infrastructure server (e.g. the membership server in NPS can act as Surveyor registrar, while in Vivaldi such server must either be introduced or at least integrated inside an existing bootstrap infrastructure). On joining the coordinate system, a node interrogates this server to obtain the identity of several (randomly chosen) Surveyors. The node then measures its distance to these Surveyors and selects the closest one as representative. From there on, the node fully complies to the embedding protocol rules, except that it will use our detection method to accept or reject embedding steps.

³ e_l is the exponential moving weighted average of the measured relative errors of all previously completed embedding steps.

However, when the node has rejected half of its current peer nodes during a same embedding round, it will seek to acquire a new filter as the high rejection rate may indicate that the filter parameters in use may have become stale (i.e. the filter needs “recalibrating”). The node then gets from its current Surveyor (or, as a fallback, any other Surveyors it knows, or the infrastructure server) the list of all the Surveyors it knows. After acquiring the current coordinates of these Surveyors, the node selects the closest one (in term of estimated distance) and obtains its Kalman filter parameters. Note that in the experiments we have carried out, which are described below, we observed very few “recalibrations”, so this very simple Surveyor selection mechanism was appropriate. However, more sophisticated approaches can be considered if need be.

5. EVALUATION

We evaluate the effectiveness of the simple anomalous/malicious behavior detection method in securing both Vivaldi and NPS. For each of these embedding protocols, we chose the most potent attack described in [11] and experimented with various populations of malicious nodes within the experimental set-up described in section 3. On PlanetLab, all these experiments were run concurrently so as to experience the same network conditions. In line with the results of section 3.3, the population of Surveyors was set to 8% of the overall population. Surveyors and malicious nodes were chosen at random.

We only present PlanetLab results, as simulation results are always equivalent or better.

5.1 Performance Metrics

To characterize the performance of our detection test, we use the classical false/true positives/negatives indicators. Specifically, a *negative* is a normal embedding step which should therefore be accepted by the test and completed. A *positive* is a malicious embedding step (i.e. where either, or both, the distance estimation and distance measurement between the node and its peer node have been tampered with) which should therefore be rejected by the test and aborted. The number of negatives (resp. positives) in the population comprising all the embedding steps is \mathcal{P}_N (resp. \mathcal{P}_P).

A *false negative* is a malicious embedding step that has been wrongly classified by the test as negative, and has therefore been wrongly completed. A *false positive* is a normal embedding step that has been wrongly rejected by the test and therefore wrongly aborted. *True positives* (resp. *true negatives*) are positives (resp. negatives) that have been correctly reported by the test and therefore have been rightly aborted (resp. completed). The number of false negatives (resp. false positives, true negatives and true positives) reported by the test is \mathcal{I}_{FN} (resp. \mathcal{I}_{FP} , \mathcal{I}_{TN} and \mathcal{I}_{TP}).

We use the notion of *false negative rate* (FNR) which is the proportion of all the malicious embedding steps that have been wrongly reported as normal by the test, and $FNR = \mathcal{I}_{FN}/\mathcal{P}_P$. The *false positive rate* (FPR) is the proportion of all the normal embedding steps that have been wrongly reported as positive by the test, so $FPR = \mathcal{I}_{FP}/\mathcal{P}_N$. Similarly, the *true positive rate* (TPR) is the proportion of malicious embedding steps that have been rightly reported as malicious by the test, and we have $TPR = \mathcal{I}_{TP}/\mathcal{P}_P$.

The *true positive test fraction* (TPTF) is the proportion of positive tests that correctly identified malicious embedding steps ($TPTF = \mathcal{I}_{TP}/(\mathcal{I}_{TP} + \mathcal{I}_{FP})$).

5.2 Securing Vivaldi

We experimented with our detection scheme on a Vivaldi system subjected to a colluding isolation attack as described in [11]. In this scenario, malicious nodes are trying to isolate a target node, by

repulsing all other nodes away from it. The malicious nodes agree on a large “exclusion” zone around the target node and randomly set their own coordinates outside this zone to try and attract honest nodes out of the exclusion zone. Note that an attacker always uses the same coordinate when lying to a given honest node.

5.2.1 Detection Method Performance

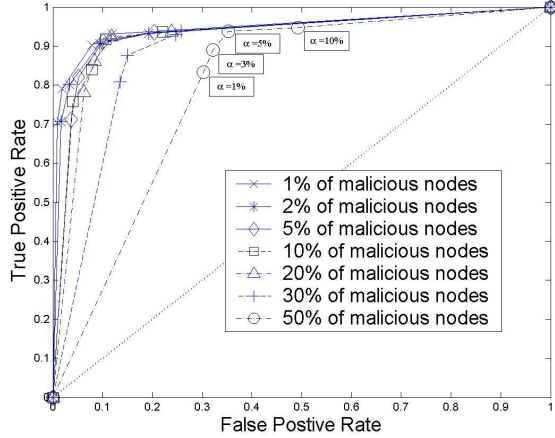


Figure 9: ROC curves. Each tick on the plots corresponds to a different value of the test’s significance level (α).

To evaluate the efficiency of the test, we first plot in figure 9, ROC (Receiver Operation Characteristics) curves observed for different significance levels (α) and several intensities of attacks. These plots show, for each significance level⁴, the point corresponding to the false positive rate along the x-axis and to the true positive rate along the y-axis, with one curve per malicious group size. Obviously, the closer to the upper left corner of the graph a curve is, the better, since such points correspond to high true positive rates (i.e. a high proportion of positives being reported as such by the test) for low false positive rates (i.e. a small proportion of negatives incorrectly reported as positives). We observe that from this perspective, the detection method can be considered to be excellent for 20% of malicious nodes or less, and still performs well even under heavy attack of up to about 30% of malicious nodes, while the power of the detection method naturally decreases as the malicious population becomes more significant. Another interesting properties of ROC curves is that they show the optimal range for the significance level. Indeed, as the slope of the ROC curve flattens, the increase in true positive rate is proportionally smaller than the corresponding increase in false positives. In other words, a higher significance level, although it always increases the true positive rate of the test, is not always productive as it eventually does more bad than good through increased false positive rates (i.e. the proportion of normal embedding steps that are aborted increases). This means that the significance level of the test should be set to a value that yields a point in the “elbow” of the ROC curve. Based on figure 9, we can deduce that a significance level of 5% seems to be a good compromise.

Figure 10 shows the true positive test fraction of the detection method for various test significance levels under various intensity of attacks. We see that the proportion of positive tests that are

⁴Significance level values α always increase as a ROC curve is “followed” from the origin. In our experiments, we used values of 1%, 3%, 5% and 10% for α .




Figure 10: True positive test fraction.

true positives is constantly high, regardless of the significance level chosen, for moderate to quite significant proportions of malicious nodes in the population (up to 20% of malicious nodes). However, thereafter the proportion of correct positive tests starts to decrease, although the rate of decrease is inversely proportional to the significance level used. This is because a higher significance level produces more positive tests, catching most malicious embedding steps, and so many more false positives are needed to make up a significant proportion of these. In light of this, a significance level of 5% offers a good compromise.




Figure 11: False Positive Rate.

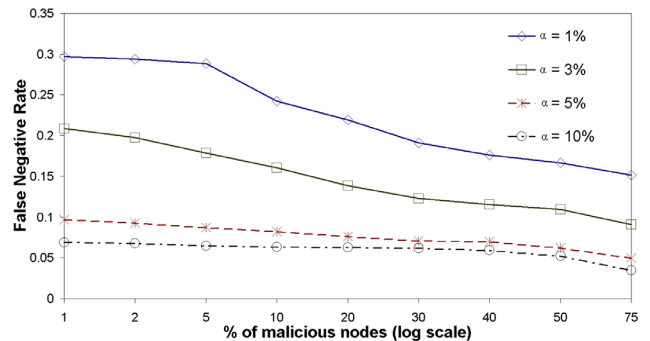


Figure 12: False negative rate.

Figures 11 and 12 show the false positive and negative rates respectively. As expected, a higher significance level results in a more aggressive test that incorrectly classifies a larger portion of

normal embedding steps (figure 11) as malicious, while a more lenient test (lower significance level) wrongly reports a higher proportion of malicious embedding steps as normal (figure 12).

Incorrect test results do have a negative impact on the embedding system: false positives artificially reduce the size of available normal nodes that can be used for normal embedding; false negatives give malicious nodes opportunities to corrupt and distort the coordinate space, which can propagate through the system and result in a greater proportion of normal nodes being identified as malicious (false positives) because of mis-positioning. This is exemplified in figure 11, where the false positive rate increases faster, as the population of malicious nodes increases, for lower values of the significance level of the test. Also, despite the fact that the false negative rate curves (figure 12) clearly exhibit negative slopes, one should note that these rates decrease much slower than the increase in malicious population. That is to say that as the number of malicious nodes in the system increases, the number of false negatives does increase, and more damage is incurred in the coordinate space. Although the accuracy of coordinate systems increases with the number of participating nodes, false negatives can therefore have a greater impact on the system than false positives and should therefore be thwarted in priority. As the false negative rates exhibited by tests with significance levels of 5% and 10% are roughly similar, while the more aggressive test yields proportionally a higher false positive rate, the significance level of 5% is a good compromise.

5.2.2 Embedding System Performance

From section 5.2.1, it should be clear that a significance level of 5% gives the overall best test performance. We therefore set the significance level to this value and assess the resistance of a Vivaldi system under various intensity of attacks.

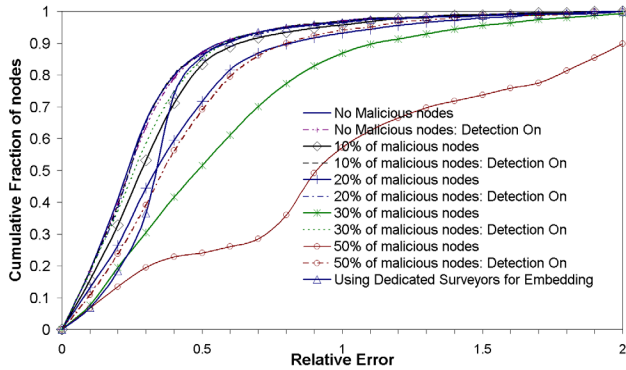


Figure 13: Distribution of measured relative errors.

The cumulative distribution function of the measured relative errors, across all normal nodes, after convergence (in the sense of error convergence as defined in [12]) is shown in figure 13 (for the time being, ignore the curve entitled "Using Dedicated Surveyors for Embedding"). We see that the detection mechanism renders the system practically immune to the attack, when the proportion of malicious nodes is 30%, or less, of the overall node population. Although the system does indeed show degraded performance for higher intensities of malicious attacks, the steeper slope of the CDF with detection, compared to the corresponding curve without (e.g. curves for 50% of malicious nodes), shows that the detection mechanism is not completely overwhelmed and still offers good protection by significantly reducing the impact of the attack.

5.3 Securing NPS

To test our proposed detection method in the context of the NPS coordinate system, we chose to study the effects of colluding isolation attack as described in [11]. The malicious nodes cooperate with each other and behave in a correct and honest way until enough of them become reference points at each layer. As soon as a minimum number of malicious reference points has been reached (in our experiments this number is set to 5) in a layer, these attackers identify a common set of victims (50% of the normal nodes they know from the layer directly below). When involved in the positioning of their victims, the malicious nodes agree to pretend they are all clustered into a remote (far away) part of the coordinate space and try and push the victims into a remote location at the "opposite" of where the attackers pretend to be, in order to isolate the victims from the other nodes in the coordinate space. In order to evade detection, including the basic detection method proposed in NPS and which is *always* turned on in our experiments, the malicious nodes use the sophisticated anti-detection method proposed in [11] during their attacks.

5.3.1 Detection Method Performance

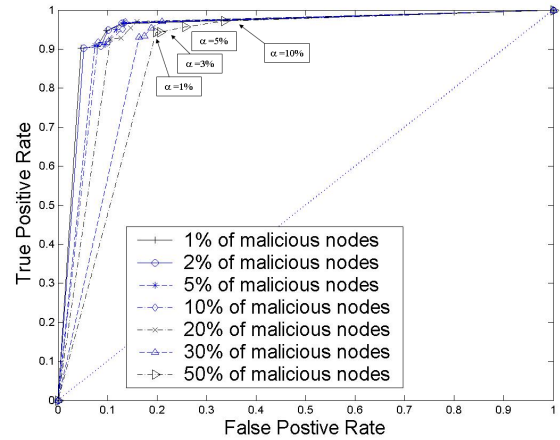


Figure 14: ROC curves.

Figure 14 shows the ROC curves for the detection test in NPS. These curves show characteristics similar to those observed in the Vivaldi system (see section 5.2.1), albeit slightly better. In particular, these curves show that the detection method withstands heavier attacks better in NPS than in Vivaldi.

There are several reasons for this. First, the basic detection method in NPS works in concert with our own, providing greater opportunities to identify malicious behavior. Also, by its very nature, the embedding method in NPS is less prone to mis-positioning error propagation amongst normal nodes, as nodes in the lower layer do not take part in the embedding of other nodes. And finally, by design, the attack considered in this section makes fewer victims than that studied in section 5.2 (i.e. 50% of normal nodes as victims vs 100% in Vivaldi).

The same observation is also true for the false positive and false negative rates (not shown) with again, overall, a significance level of 5% seemingly offering the best compromise between "catching" malicious embedding steps while not being overly cautious and over-reacting to normal variations in network conditions.

The similarities between the test performance under NPS and Vivaldi, despite the different nature of the attacks under consideration

and even differences in coordinate “structure” (two-dimensional with height for Vivaldi versus eight-dimensional for NPS), illustrates the generality of the proposed detection method. This is because our detection test is based on the modeling of a dimensionless quantity (the relative error) which is at the very core of *any* embedding method.

5.3.2 Embedding System Performance

We study the performance of the NPS embedding system when subject to increasing intensity of attacks, while being protected by our detection scheme. Note that in this section, “detection off” really means that our proposed detection mechanism is not used, but the basic NPS detection mechanism is still “on”.




Figure 15: Distribution of measured relative errors.

Figure 15 shows the cumulative distribution function of relative errors in the system. We note again similarities with the dynamic behavior of similar Vivaldi systems, except that the tail of the CDF for 50% malicious nodes with detection is heavier than the corresponding curve in the Vivaldi case. Keeping in mind that in NPS not all nodes are victims and that not all normal nodes will propagate mis-positioning errors, this indicates that the attack is still quite effective against its victims, albeit “dampened” by the detection mechanism. This effect is compounded by the fact that, with our simple detection protocol, malicious nodes that have found their way into the layer hierarchy of NPS and act as Reference Points, do stay in place throughout the experiment, despite numerous detections of their corrupt embedding steps.

Nevertheless, the detection method proposed affords near immunity to the system up to rather severe attack conditions (e.g. about 30% of malicious nodes in the system).

6. DISCUSSIONS AND CONCLUSIONS

We have presented a method for malicious behavior detection to secure the embedding phase of Internet coordinate systems. Our method does not rely on the geometric properties of the coordinate space, and is therefore unaffected by potential triangular inequality violations which often occur in the Internet [21, 22]. Instead, our detection test is based on the modeling of the dynamic relative errors observed in a clean system. The relative error is a dimensionless quantity which is at the very core of *any* embedding method, leading us to believe that our proposed detection test can effectively identify malicious behavior in very many embedding protocols and coordinate space structures that are under a potential very large range of attacks. The experiments presented in this paper do show that the performance of the detection test is effectively the same in two different scenarios involving different embedding protocols and different attacks. As far as we know, this is the first such

general detection test, capable of surviving sophisticated attacks. Also, we consider exclusively attacks aimed at distorting the coordinate space, carried out by nodes inside the embedding system. Our method thus succeeds where more obvious methods based on authentication would fail.

In practice, we introduced the concept of Surveyor nodes which, by design, are immune to embedding attacks and do observe the properties of the coordinate system in clean conditions. The Surveyors make up the basis of a “security infrastructure”. It is important to note that the deployment of Surveyors does not equate to imposing an embedding infrastructure: peer-to-peer based embedding systems, like Vivaldi, do retain their infrastructure-less embedding characteristics. Indeed, apart from a test to accept or filter out embedding steps, our method does not entail any change to the operations of the embedding protocols. Even though this paper does not address the problem of external attacks on the infrastructure (e.g. denial of service attacks, “link clogging”, etc.), we note that solutions to such attacks have been proposed elsewhere (e.g. [24]).

The operations of the proposed detection protocol were deliberately kept simple and tested on systems where Surveyors were chosen randomly, although their representativeness increases with closeness to their “clients”. Despite the possibly non-optimal Surveyor distribution resulting from such a choice, the results obtained show the effectiveness of our proposal in securing Internet coordinate systems. Nevertheless, given the enhanced coordinate service afforded by our detection test and simple detection protocol, one might envisage that ISPs may readily want to deploy Surveyors within their network to offer enhanced coordinate service to their customers. Such business-driven strategic deployment can only improve representativeness of Surveyors, and thus improve the security of large-scale coordinate service, with possibly a much smaller proportion of Surveyors than the upper bound reported in this paper (as illustrated by the simple k-means deployment in section 3.3). More sophisticated Surveyor selection mechanisms than those presented in this paper could also result in better security through better representativeness in large-scale coordinate systems. We believe that the increased robustness provided by our proposal could act as a catalyzer to the acceptance and deployment of large-scale coordinate services in the Internet.

Nevertheless, with a trusted Surveyor infrastructure in place, it could be argued that using these Surveyors for positioning other nodes would ensure immunity to any insider attacks. For Vivaldi, using the Surveyors for positioning would mean that normal nodes only choose Surveyors as neighbors. The embedding performance of such Vivaldi scenario is depicted by the “using dedicated Surveyors for embedding”-curve of figure 13, with the 1% Surveyors of the simple k-means deployment method. It can clearly be seen that such use of the Surveyor infrastructure trades embedding accuracy for increased security. A similar NPS scenario, where only Surveyors would be chosen as Landmarks and reference points, unsurprisingly led to embedding performance equivalent to a clean NPS system, as the hierarchical embedding structure in both systems are very similar. An NPS system where only Surveyors would be chosen as Landmarks and reference points actually looks like a hybrid system between GNP [1] and NPS: it is a fixed infrastructure system (like GNP), but with distributed Landmark coordinate computation and a hierarchical structure (like NPS). In both cases above, a clear scalability issue arises as the load on each Surveyor increases as their number decreases. In light of the discussion on strategic Surveyor deployment, as well as lower bound on the number of required Surveyors (see above and section 3.3), it is not clear that the solution of embedding against Surveyors only is practically viable. Even if it were, a hybrid solution, where Surveyors would be

used for malicious activity detection under mild to medium attack intensities and exclusively used for embedding under more severe conditions, would be more accurate and afford better scalability.

Furthermore, each node using malicious behavior detection does so in isolation, as there is no cooperation between nodes in a bid to improve detection and identification of malicious nodes. Instead, any embedding step identified as malicious by a node is simply, and quietly, ignored and discarded locally. One of the reasons the detection protocol was designed in this way was to avoid potential denial-of-service attacks that could result from the sharing of information about malicious activity, with the view of excluding offending nodes. Indeed, such an approach could open the door to an attack that consists in trying to get honest nodes excluded from the system through the collusion of wrong malicious reports. Trust propagation could be used to mitigate or remove this threat and thus allow detection cooperation amongst nodes, which could only improve the security of the overall system and push the boundary of applicability of the detection by reducing further the impact of very large-scale attacks on the embedding system.

A property of coordinate systems that has been observed in large-scale deployment is coordinate drift [7]. While the drift rate is low enough that it does not interfere with our detection method, and can thus be ignored, drift, and in general coordinate variations, can nevertheless be an issue for applications using the coordinate system for distance estimation (the “usage phase” of the coordinate service). Indeed, even though the embedding phase of the system may have been secured, this would not prevent a malicious node from blatantly lying about its coordinates when a node requests them for simple distance estimation during normal use of the service. This normal use of the service must therefore also be secured, perhaps through the use of validity period for “certified” coordinates. This is left as future work. However, because lying about its own coordinate during a simple distance estimation only fools a single other node at a time, this type of application-level attack is less dangerous than an attack on the embedding phase itself which can distort the coordinate space and thus spoil the coordinate computations of unsuspecting honest nodes (and thus bias distance estimates involving these nodes). This is the reason why we chose to address the problem of securing the embedding phase first.

Acknowledgments

We would like to thank Dina Papagiannaki, Guy Leduc and the anonymous reviewers for their useful comments that helped improve the paper.

7. 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 (IPTPS03), 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, New York, June 2002.
- [7] J. Ledlie, P. Gardner, and M. Seltzer, *Network Coordinates in the Wild*, in Proceedings of NSDI, Cambridge, MA, April 2007.
- [8] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, *Chord: A Scalable Peer-to-peer Lookup Service for Internet Applications*, in Proceedings of SIGCOMM, San Diego, CA, August 2001.
- [9] Azureus BitTorrent Client. <http://azureus.sourceforge.net>
- [10] www.skype.com
- [11] M. A. Kaafar, L. Mathy, T. Turletti, and W. Dabbous, *Virtual Networks under Attack: Disrupting Internet Coordinate Systems*, in Proceedings of CoNext 2006, Lisboa, December, 2006.
- [12] 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),Pisa, September 2006.
- [13] R.E. Kalman, *A New Approach to Linear Filtering and Prediction Problems*, in Transactions of the ASME - Journal of Basic Engineering Vol. 82: pp. 35-45, 1960.
- [14] R.E. Kalman, and R.S. Bucy, *New Results in Linear Filtering and Prediction Theory*, in Transactions of the ASME - Journal of Basic Engineering Vol. 83: pp. 95-107, 1961.
- [15] 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, CA, November 2001.
- [16] Z. Ghahramani, G. Hinton, *Parameter Estimation for Linear Dynamical Systems*, University of Toronto, Technical Report CRG-TR-96-2.
- [17] M. A. Kaafar, L. Mathy, C. Barakat, K. Salamatian, T. Turletti and W. Dabbous, *Securing Internet Coordinate System: Embedding Phase*, Technical Report INRIA inria-00151257.
- [18] K. P. Gummadi, S. Saroiu, and S. D. Gribble, *King: Estimating Latency between Arbitrary Internet End Hosts*, in Proceedings of SIGCOMM Internet Measurement Workshop (IMW), Pittsburgh, PA, November 2002.
- [19] H. Lilliefors, *On the Kolmogorov-Smirnov test for normality with mean and variance unknown*, Journal of the American Statistical Association, Vol. 62. pp. 399-402, June, 1967.
- [20] A. Soule, K. Salamatian, and N. Taft, *Combining Filtering and Statistical Methods for Anomaly Detection*, in Proceedings of Internet Measurement Conference (IMC), Berkeley, October, 2005.
- [21] 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.
- [22] 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.
- [23] J. Bilmes, *A gentle tutorial on the EM algorithm including gaussian mixtures and baum-welch*, Technical Report TR-97-021, International Computer Science Institute, Berkeley, CA, 1997.
- [24] A. Keromytis, V. Misra and D. Rubenstein, *SOS: Secure Overlay Services*, in Proceedings of ACM SIGCOMM, Pittsburgh, PA, August 2002.