Alcom-FT Technical Report Series
ALCOMFT-TR-03-153

Attack Propagation in Networks*

Sotiris Nikoletseas† Grigorios Prasinos† Paul Spirakis† Christos Zaroliagis†‡

August 1, 2002

Abstract

A new model for intrusion and its propagation through various attack schemes in networks is considered. The model is characterized by the number of network nodes n, and two parameters f and g. Parameter f represents the probability of failure of an attack to a node and is a gross measure of the level of security of the attacked system and perhaps of the intruder's skills; g represents a limit on the number of attacks that the intrusion software can ever try, due to the danger to be discovered, when it issues them from a particular (broken) network node. The success of the attack scheme is characterized by two factors: the number of nodes captured (the spread factor) and the number of virtual links that a defense mechanism has to trace from any node where the attack is active to the origin of the intrusion (the traceability factor). The goal of an intruder is to maximize both factors. In our model, we present four different ways (attack schemes) by which an intruder can organize his attacks. Using analytic and experimental methods, we first show that for any $0 < f < 1$, there exists a constant g for which any of our attack schemes can achieve a $\Theta(n)$ spread and traceability factor with high probability, given sufficient propagation time. We also show for three of our attack schemes that the spread and the traceability factors are, with high probability, linearly related during the whole duration of the attack propagation. This implies that it will not be easy for a detection mechanism to trace the origin of the intrusion, since it will have to trace a number of links proportional to the nodes captured.

1 Introduction

Attacks in computer networks pose several key problems regarding intrusion propagation and detection [4, 8, 18, 19]. Various models have been proposed under which researchers mainly study the effective detection and defeat of attacks assuming a very powerful intruder; see e.g., [15]. In this setting, intrusion propagation (the process of spread of such attacks) has mostly been investigated under gossip or epidemiological models [9, 10, 11]. On the other hand, the fear of malicious attacks along with the development of advanced cryptographic techniques has considerably increased the security level of current computer systems. Contrary to previous studies, we are interested here in studying intrusion propagation assuming that the intruder has a rather limited power and would like to investigate how intrusion can propagate in a perhaps highly secure network. To this end, we introduce a general model for such an intrusion and its propagation in networks.

Let $\mathcal{N}$ be an n-node network, e.g., a computer network, with some physical infrastructure underlying it and whose specific topology is not a concern. By the term "network", we do not

*This work was partially supported by the Future and Emerging Technologies Programme of EU under contract no. IST-1999-14186 (ALCOM-FT), by the Human Potential Programme of EU under contracts no. HPRN-CT-1999-00104 (AMORE) and HPRN-CT-1999-00112 (ARACNE), and by the Greek GRST project ALKAD.

†Computer Technology Institute and Department of Computer Engineering & Informatics, University of Patras, 26500 Patras, Greece. E-mails: nikole@cti.gr, green@ceid.upatras.gr, spirakis@cti.gr, zaro@ceid.upatras.gr.

‡Corresponding author.

necessarily mean that communication is done point-to-point. We rather view a network as a collection of host systems each one having its own logical address (e.g., the Internet), and whose underlying physical infrastructure uses advanced communication and interconnection technology such as buses, optical links, wireless communication media, etc. Direct communication between two nodes of $\mathcal{N}$ is achieved by establishing a *virtual* channel through the physical infrastructure between these two nodes.

Assume that in such a network an intruder, starting from his own computer, would like to break as many other systems as possible. The intrusion consists of a collection of attacks. An *attack* is issued from some node in $\mathcal{N}$ and is an attempt to break the perimetric security of another node (host system) in $\mathcal{N}$. The intrusion is realized by an attack scheme. An *attack scheme* is a protocol for the organization of the attacks issued from specific nodes of $\mathcal{N}$. The intruder is a greedy one, i.e., does not have a specific target, and attacks computer systems equiprobably at random[1]. An attack succeeds or fails, independently of other attacks, with a failure probability $0 < f < 1$ that represents the difficulty of breaking a system in $\mathcal{N}$; f is a gross measure of the security level of the attacked systems (e.g., of the average security or the perceived maximum security level of a system) and may also depend on the intruder's skills. Our model assumption about f is motivated by a large class of existing attacks; for example, attacks that are based on randomly sampling a set of possible passwords from a large password domain and then trying each of them. Many practical local attack software programs (e.g., a cracking password worm) work in this way. Furthermore, the probability of success of such a scheme in a node does not depend on previous successes at other nodes or on previous attempts at the same node. This is because the locally implemented set of passwords is perhaps different in each node and the set of passwords used by the local attack software is very small (for reasons of speed) compared to the password domain set.

If an attack does not fail, then some, randomly and equiprobably chosen, network node is returned. Because of that, it may happen that an already selected node (an already broken system) is chosen again. If the result of a non-failed attack is a node which has not been selected before, then the attack is considered *successful* and a *virtual link* (virtual channel) is established to that node. The random selection, with possible repetition, of a node in the case of a non-failed attack is motivated by the following pragmatic considerations: (i) if the local attack software (e.g., a worm) is successfully confronted, it should not reveal any information about broken nodes in the past; (ii) the local attack software may blindly extend attacks to hosts contained in tables of the newly broken systems which may include the already broken ones.

The intruder tries to protect himself as much as possible from being traced: once a system is broken, his software tries from *that* system to attack (again equiprobably at random) another system by disguising itself as a user process of the broken system. Because of the danger of being discovered, the intruder's software can only ever try a limited (i.e., constant) number g of attacks from a specific node of the network. If a successful attack is issued before the limit g is reached, then the software enters a dormant phase and performs no action (in order to not raise any suspicions). If at some node i the software exhausts the attack bound g, then it terminates execution at i and "backtracks" to a previously broken system j to continue from there its attacks, provided that there are still some attempts left at j. In such a case, the local software at j is reactivated and starts again to issue attacks. If at any time during the execution of the attack scheme, the intrusion is discovered by some system, we assume that the whole attack scheme to $\mathcal{N}$ terminates.

Two natural questions raised here are: (a) how long the intruder can go, i.e., how many computer systems can be successfully attacked in $\mathcal{N}$ until he is discovered, and (b) how many virtual links a detection mechanism has to trace in order to find the origin of the intrusion. In particular,

[1]One way to accomplish this is to choose a random IP address from the space of all possible IP addresses; actually, the intruder needs only to know the valid ranges of each field of an IP address and then choose for each field a random number within its valid range.

assume that the intrusion starts at time 0 with attack scheme S. At any time $t \geq 0$, let $n_S(t)$ be the number of nodes captured, called the *spread factor*, and let $\ell_S(t)$ be the shortest distance (in number of virtual links) from the currently active position of the intruder's software to the origin, called the *traceability factor*. Given a discovery (i.e., stopping) time T, we would like to estimate $n_S(T)$ and $\ell_S(T)$. We shall refer to this as the *attack propagation* problem. The goal, from the side of the intruder, is to employ an attack scheme which maximizes *both* factors. It is not at all obvious how an intrusion can be organized so that both $n_S(T)$ and $\ell_S(T)$ are large. In fact, to the best of our knowledge, almost all epidemiological (and gossip propagation) models have usually very small $\ell_S(T)$ compared to $n_S(T)$, because of their "radially" spreading nature.

The above process defines naturally a graph G whose vertices correspond to the nodes of the network and if a virtual link (i.e., a successful attack through some virtual channel) is established between two nodes i and j, then an edge between vertices i and j is added to G. In this setting, the spread factor $n_S(T)$ is the size of the obtained connected component in G, and the traceability factor $\ell_S(T)$ is the length of the shortest (in number of edges) path in this connected component from the current vertex (node) issuing attacks to the origin (the node from which the intrusion was started). We shall refer to this path as the *traceability path*. Hence, the attack propagation problem reduces in estimating the values of these two quantities in G.

Our work is centered around the attack propagation problem. We present four different attack schemes (protocols) by which an intruder can organize his attacks in the above model. Our starting point is an attack propagation scheme that organizes attacks along a single traceability path. This scheme, inspired by ideas developed in [13] for a different problem and setting, forms the basis for the development of three other attack schemes, called tree attack schemes (or tree protocols). Using analytic and experimental methods, we first show that for *any* $0 < f < 1$, there exists a g for which any of our attack schemes will achieve a $\Theta(n)$ spread factor with high probability, provided T is sufficiently large. This means that if an intrusion is realized by any of our attack schemes, it will spread, regardless of the security level, to a big part of the network. We also show that the spread and the traceability factors are linearly related. Actually, for our tree attack schemes this linear relationship holds, with high probability, during the *whole* duration of the attack propagation. This implies that it will not be easy for a detection mechanism to trace the origin of the intruder, since at any time it will have to trace a number of links proportional to the number of nodes captured.

Another interesting issue is to investigate the possibility of a total failure of such attack schemes, namely the possibility that they eventually return to their starting point, not because the intruder is discovered but due to backtracking caused by the limited number of attempts from a specific node. Our combined analytic and experimental methods show that for all four attack schemes such a possibility is very small.

It is worth mentioning that our analytic and experimental methods are tied to each other. The analytic study of the protocols, where applicable, is rather complicated and gives only lower bounds that are (probably) not tight. Hence, we have to resort to experiments to get insight as well as a basis of reasonable assumptions to further proceed with the analysis.

The remainder of the paper is organized as follows. In Section 2, we describe in detail the model of intrusion and its propagation that we consider and provide a comparison with other models. Our attack schemes and their analysis is presented in Section 3, while the results of their experimental evaluation are given in Section 4. Finally, in Section 5 we conclude and argue that our work spawns several interesting research issues that require further investigation. A preliminary version of this work appeared in [14]. Since some of the results originate from ideas in [13], the current version constitutes an integration of the results in [13, 14].

2 The Model

In this section, we shall describe in more detail the model we consider in this work.

We assume that we are given a network $\mathcal{N}$ consisting of a set V of n nodes. We further assume that some physical network infrastructure exists underlying $\mathcal{N}$ and its specific topology is not a concern of the model. The nodes of $\mathcal{N}$ can be in two states: *awake* and *sleeping*. An awake node may (or may not) possess a special token. Awake nodes possessing a token are called *active*. Only active nodes are allowed to perform attacks. An *attack* is an attempt from an active node to break the security of some other node in $\mathcal{N}$. The active nodes constitute the "frontier" of the awake nodes which issue attacks. Initially, exactly one node x_0 is awake and active (representing the original position of the intruder), and the rest are considered sleeping. Each node is allowed a maximum, but fixed, number g of attacks for establishing a virtual link with some other node of $\mathcal{N}$. Attacks from each active node are executed one-by-one and each may independently fail with probability $0 < f < 1$. If an attack does not fail, then a randomly chosen node v from V is returned. If v is sleeping, then a virtual link to v (through some virtual network channel) is established and the attack is considered *successful*. That is, established links represent successful attacks. Once a link is established, the sending node is notified and the receiving node (i.e., v) becomes awake. The model allows repetitions, that is, the same node may be returned in two different attacks which did not fail (but obviously in such a case the attack is not considered successful).

Our model can be naturally represented by a graph G whose vertex set is V and if a virtual link (successful attack) is established between two nodes i and j of $\mathcal{N}$, then an edge between vertices i and j is added to G.

An *attack scheme* in such a model is a distributed computation protocol which specifies how the active nodes are created as the intrusion proceeds (the terms "attack scheme" and "protocol" will be used interchangeably throughout the paper). The protocol also specifies information exchanges between active and awake nodes. Clearly, such information exchanges are carried out along the established links (edges of G). It is a duty of the protocol to maintain information about the currently established links.

Other models that study intrusion propagation are, to the best of our knowledge, mostly based on approaches from biological epidemiology (see e.g., [9, 10, 15]). These models are analyzed with various stochastic processes varying from relatively simple [9, 10] to more sophisticated approaches – an example of the latter is the use of the theory of interacting particle systems through the model of contact processes [11]. Another modeling could be achieved through the use of dynamic games in networks, in particular of dynamic monopolies (dynamos) [16, 17], which is somehow similar to the epidemiological models. All these models usually assume a very powerful intruder and they mainly watch the (very rapid) evolution of the intrusion (viruses, worms, etc) in terms of its birth/death rates and the topology of connections between systems (non-infected nodes become infected at a rate proportional to the number of infected neighbors). In the case of dynamos, almost all research has focused on the study of patterns of initially infected nodes whose occurrence could lead the entire network to total infection [6, 7, 16, 17].

The above models, however, do not provide any kind of parameterization of the power of the intruder, or any algorithmic strategy for the spread of the intrusion. In addition, as noted in [21], simplified epidemiological models may fail to model adequately the spread of the intrusion. In contrast, our model allows for the characterization (and/or control) of the power of the intruder with the parameters f and g (the smaller the f, the more powerful the intruder, and vice versa for g), and also allows the intruder to develop strategies in order to spread the intrusion. For these reasons, we feel that our model is closer to a more realistic modeling of intrusion propagation. In Section 5, we discuss further extensions of our model.

3 The Attack Schemes and their Analysis

A protocol for attack propagation in a network $\mathcal{N}$ according to our model can be derived by extending ideas developed in [13]. In particular, attack propagation can be achieved by a suitable attack scheme which incrementally (link by link) extends and maintains a traceability path of attacked systems using a protocol that comes in two versions (for the sake of technical analysis). The first version assumes the existence of a special capability that restarts the protocol in the case of a total failure and which always succeeds. The second version rectifies this strong assumption by basically simulating this special capability and hence making it unnecessary. In our setting, we shall refer to these two versions as attack schemes or protocols S_1 and S_2 respectively.

We will assume for S_1 that there exists a so-called *special attack* $A(x, U)$ which can be issued by any active node x to a subset $U \subseteq V$ of sleeping nodes and that always succeeds. After $A(x, U)$ has been issued, a node x' of U is selected randomly, is made active, and the protocol restarts execution in the subnetwork consisting of nodes in $U - \{x'\}$. Later we shall show how S_2 simulates this special attack and makes it unnecessary.

Both protocols assume that there is only one active node, i.e., there is only a single token possessed by some awake node, representing the current position from which attacks are issued. The particular awake node possessing the token is determined by the protocol. Protocols S_1 and S_2 are the starting points of our study. Based on these, we develop three other variants. All protocols are described in the remainder of this section.

3.1 The basic attack scheme S_1

In attack scheme S_1 the g attacks per node are grouped into two equally sized sets, called the *green* and the *red* set, respectively; that is, $g = 2\lambda$ where λ is the cardinality of each set (green or red). The separation of the attacks into two sets is done only for the sake of technical analysis, to provide stochastic independence. The main idea of S_1 is to organize the attacks along the traceability path of attacked systems. The protocol tries initially to establish a (long) traceability path, link by link, using only the green attacks. Each attack is issued from the last node in the path which is the active node (i.e., it possesses the token). An attack is considered successful if a sleeping node is returned which will now become the new last node of the path and gets the token. When attack propagation, i.e., extension of the constructed path, using green attacks is not possible, then the red set of attacks is used. If extension to a new node is established, then S_1 passes the token to that node and continues from there using its green attacks. Otherwise, S_1 backtracks to the first node whose red attacks have not been used yet and (after passing the token) tries to extend the path from that node using the red attacks.

Less informally, S_1 performs a number of logical steps. At the end of the k-th logical step the protocol maintains a triple (Π_k, U_k, R_k), where Π_k is the traceability path constructed so far, U_k is the set of sleeping nodes, and R_k is the set of *red nodes*, i.e., of those awake nodes whose green attacks have all been used. Initially $k = 0$, $\Pi_0 = \{x_0\}$ (the first active node), $U_0 = V - \{x_0\}$, and $R_0 = \emptyset$. Let x_k (resp. z_k) be the first (resp. last) node of Π_k; z_k is assumed to hold the special token denoting the node issuing attacks. S_1 performs the following logical step until $U_k = \emptyset$ or all nodes have exhausted their attacks. There are two cases to consider depending on whether z_k belongs to R_k or not.

Case 1: $z_k \notin R_k$. Node z_k tries its green attacks one-by-one. Every such attack either fails or a random node v is returned. If $v \in U_k$, then the traceability path is extended by setting $z_{k+1} = v$, $\Pi_{k+1} = \Pi_k \cup \{z_{k+1}\}$, $U_{k+1} = U_k - \{z_{k+1}\}$, $R_{k+1} = R_k$, and the token is passed to z_{k+1}. Otherwise, $\Pi_{k+1} = \Pi_k$, $U_{k+1} = U_k$, $R_{k+1} = R_k \cup \{z_k\}$.

Case 2: $z_k \in R_k$. Node z_k tries its red attacks one-by-one. Every such attack either fails or a random node v is returned. If $v \in U_k$, then the traceability path is extended by setting $z_{k+1} = v$,

$\Pi_{k+1} = \Pi_k \cup \{z_{k+1}\}$, $U_{k+1} = U_k - \{z_{k+1}\}$, $R_{k+1} = R_k$, and the token is passed to z_{k+1}. Otherwise, z_k sends a message backwards in Π_k to find the first node whose red attacks have not been tried yet. Let us call this node z_{k+1}; this node also receives the token from z_k. Now, Π_{k+1} is the subpath of Π_k from x_k up to z_{k+1}, $U_{k+1} = U_k$, and $R_{k+1} = R_k \cup \{z_{k+1}\}$. In the special case where $z_{k+1} = x_k$ (i.e., when the traceability path has shrunk to a single node) and all of its red attacks have been unsuccessfully tried, then z_{k+1} sends the special attack $A(z_{k+1}, U_k)$. Subsequently, S_1 starts its execution from a new active node x_{k+1}, i.e., $\Pi_{k+1} = \{x_{k+1}\}$, $U_{k+1} = U_k - \{x_{k+1}\}$, and $R_{k+1} = R_k$.

The attack scheme S_1 can be easily implemented in a distributed way. To implicitly maintain the triple (Π_k, U_k, R_k), each node maintains two pointers, one to its current successor node and one to its predecessor node in the traceability path, a flag indicating whether it is a red node, and a counter indicating how many attacks have still been left.

3.2 Structural properties of S_1

We start with the structural properties of S_1, as it will form the basis of our analysis regarding the rest of the attack schemes. We shall show that S_1 can achieve a $\Theta(n)$ spread and traceability factor and that the expected number of special attacks for restarting purposes is constant. In the following, we shall say that a statement Q holds *with high probability*, if $\Pr[Q] = 1 - o(1)$ as $n \to \infty$.

Let $\ell_k = |\Pi_k| - 1$ (length of Π_k), $u_k = |U_k|$, and $r_k = |R_k|$. Clearly, $\ell_0 = 0$, $u_0 = n - 1$, and $r_0 = 0$, which implies that $n - u_0 + r_0 = 1$. Notice that for S_1, we have $n_{S_1}(k) = n - u_k$ and $\ell_{S_1}(k) = \ell_k$. In each logical step k, S_1 increases either $n - u_k$ or r_k by at most one. Consequently, $(n - u_k) + r_k \leq (n - u_{k-1}) + r_{k-1} + 1$ and inductively we get

$$(n - u_k) + r_k \leq k + 1. \tag{1}$$

From the description of the protocol, it is clear that $V - U_k - \Pi_k \subseteq R_k$, which in turn implies that $n - u_k - |\Pi_k| \leq r_k$. Since $\ell_k = |\Pi_k| - 1$, we get that $\ell_k \geq n - u_k - r_k - 1 \geq n - u_k - (k + 1 - (n - u_k)) - 1$, where the second inequality follows by Eq. (1). Hence,

$$\ell_k \geq 2(n - u_k) - k - 2. \tag{2}$$

Let $\mathcal{A}_i$ be the set of attacks issued at step i. Then, $H_k = \{\Pi_i, U_i, R_i, \mathcal{A}_i\}_{i=0}^{k}$ represents the history of the protocol up to step k. The probability of failing to extend the path is

$$\begin{aligned} \Pr[u_{k+1} = j | u_k = j, H_k = H] &= \sum_{t=0}^{\lambda} \binom{\lambda}{t} \left(\frac{f u_k}{n}\right)^t \left(1 - \frac{u_k}{n}\right)^{\lambda - t} \\ &= \left(1 - \frac{u_k}{n} + \frac{f u_k}{n}\right)^{\lambda} = \left(1 - \frac{u_k}{n}(1 - f)\right)^{\lambda} \end{aligned} \tag{3}$$

for all k, j and H. Since this probability depends only on u_k, and since u_{k+1} is either u_k or $u_k - 1$, it follows that the sequence $\{u_k\}$ is a Markov chain. By considering the sojourn times of u_k in each state of the chain, we can prove the following.

Theorem 1 *Let $\lambda = g/2$, $\lambda(1 - f) \geq 4 \ln 2$ and $r = \lambda(1 - f)/n$. Then, for attack scheme S_1 there exists a time $t_0 = n - e^{-rn}/r + o(n)$ such that if $T \geq t_0$, then $n_{S_1}(T) \geq n(1 - 2\ln 2/g(1 - f))$ and $\ell_{S_1}(T) \geq n(1 - 4\ln 2/g(1 - f))$ with high probability.*

Proof. Recall that $n_{S_1}(k) = n - u_k$ and $\ell_{S_1}(k) = \ell_k \geq 2(n - u_k) - k - 2$ (Eq. (2)). Hence, providing lower bound estimations for the spread and traceability factors boils down to providing an upper bound j for u_k as well as to showing that $2j + k$ is small (as required by Eq. (2)).

We shall use the fact that the sequence $\{u_k\}$ is a Markov Chain. Let $X_i = \max\{|k-l| : u_k = u_l = i\}$ for $1 \le i \le n-1$. Clearly, $X_i + 1$ is the length of the *sojourn time* of u_k in state i, and let $Y_j = \sum_{i=j+1}^{n-1}(X_i+1)$, i.e., the *hitting time* of state j of u_k. By definition, it can be easily proved that

$$\Pr[u_k \le j] = \Pr[Y_j \le k]. \tag{4}$$

Hence, it suffices to exhibit a pair (j,k) for which $\Pr[Y_j \le k] = 1 - o(1)$ and $2j+k$ is small.

The fact that $\{u_k\}$ is a Markov chain implies that $\{X_i\}$ are *independent* geometric random variables. By Eq. (3), their success probability is $p = 1-(1-\frac{i}{n}(1-f))^\lambda$. The mean and variance of $\{X_i\}$ are given by $E[X_i] = (1-p)/p$ and $\sigma^2[X_i] = (1-p)/p^2$. Let $r = \frac{\lambda}{n}(1-f)$. Since $(1-\frac{i}{n}(1-f))^\lambda \le e^{-ri}$, we have that

$$E[X_i] \le \frac{e^{-ri}}{1-e^{-ri}} \qquad \text{and} \qquad \sigma^2[X_i] \le \frac{e^{-ri}}{(1-e^{-ri})^2}.$$

By choosing $j = \left\lceil \frac{\ln 2}{r} \right\rceil$ we get that

$$\begin{aligned}
\sigma^2(Y_j) &= \sum_{i=j+1}^{n-1} \sigma^2[X_i+1] = \sum_{i=j+1}^{n-1} \sigma^2[X_i] \\
&\le \sum_{i=j+1}^{n-1} \frac{e^{-ri}}{(1-e^{-ri})^2} \le \sum_{i=j+1}^{n-1} \frac{e^{-rj}}{(1-e^{-rj})^2} \\
&\le \sum_{i=j+1}^{n-1} \frac{1/2}{(1-1/2)^2} \le 2(n-j-1) \\
&< 2n.
\end{aligned}$$

We also need to upper bound $E(Y_j)$.

$$\begin{aligned}
E[Y_j] &= \sum_{i=j+1}^{n-1} E[X_i+1] \le \sum_{i=j+1}^{n-1} \frac{1}{1-e^{-ri}} \\
&\le \int_j^{n-1} (1-e^{-rx})^{-1}\,dx \le \frac{1}{r}\int_{\ln 2}^{rn} (1-e^{-y})^{-1}\,dy \\
&= \frac{1}{r}\int_{\ln 2}^{rn} \left(1 + \sum_{m=1}^{\infty} e^{-my}\right) dy \qquad \text{(Taylor series expansion)} \\
&= \frac{1}{r}\left[y - \sum_{m=1}^{\infty} \frac{e^{-my}}{m}\right]_{\ln 2}^{rn} = n - \frac{\ln 2}{r} - \frac{1}{r}\sum_{m=1}^{\infty}\frac{e^{-mrn}}{m} + \frac{1}{r}\sum_{m=1}^{\infty}\frac{2^{-m}}{m} \\
&< n - \frac{\ln 2}{r} - \frac{e^{-rn}}{r} + \frac{1}{r}\ln 2 \\
&= n - \frac{e^{-rn}}{r}.
\end{aligned}$$

Thus, by applying the following version of Chebyshev's inequality

$$\Pr[Y_j < E(Y_j) + h(n)\sigma(Y_j)] \ge 1 - \frac{1}{h(n)}$$

and by choosing any function $h(n) = o(\sqrt{n})$ we get

$$\Pr\left[Y_j < n - \frac{e^{-rn}}{r} + h(n)\sqrt{n}\right] = 1 - o(1)$$

as $n \to \infty$. Hence, with high probability, and for any time $T \geq t_0 = n - \frac{e^{-rn}}{r} + h(n)\sqrt{n}$, we have by Eq. (2) and (4) that

$$n_{S_1}(T) \geq n - u_k \geq n - \frac{\ln 2}{r} = n\left(1 - \frac{\ln 2}{\lambda(1-f)}\right) = n\left(1 - \frac{2\ln 2}{g(1-f)}\right)$$

$$\ell_{S_1}(T) \geq 2\left(n - \frac{\ln 2}{r}\right) - n + \frac{e^{-rn}}{r} - h(n)\sqrt{n} - 2 \geq n\left(1 - \frac{2\ln 2}{\lambda(1-f)}\right) = n\left(1 - \frac{4\ln 2}{g(1-f)}\right).$$

∎

We now turn to the estimation of the number of special attacks required by protocol S_1 for restarting purposes. We shall show that the stochastic process ℓ_k is dominated by a biased random walk in one dimension. The required result will follow by upper bounding the expected number of returns of the random walk to the origin.

Let p be the probability that ℓ_k is extended by 1 via green attacks, and let $q = 1 - p$ be the probability that ℓ_k fails to extend and shrinks to a previous node of unused red attacks. Then, by Eq. (3) $p = 1 - (1 - \frac{u_k}{n}(1-f))^\lambda$.

Let α be a constant. Consider the first $\tau_0 = \alpha \log n$ green attacks of S_1, and let I_0 be the time interval of these attacks. Then, the length of the traceability path achieved during I_0 is $\Omega(\log n)$ with high probability, as the following lemma shows.

Lemma 1 *During the time interval I_0, the established traceability path has length $\mu_0 \geq (1-\gamma)\frac{3}{4}\alpha \log n$, for any $\gamma \in (0,1)$ and $\alpha > \frac{8}{3\gamma^2}$, with probability at least $1 - n^{-d}$, $d > 1$.*

Proof. Since $u_0 = n - 1$ it follows, even if all green attacks were successful, that $u_{\tau_0} \geq n - 1 - \alpha \log n \geq \beta n$, for some $\beta < 1$. For the time interval I_0, we then have

$$\begin{aligned} p &= 1 - \left(1 - \frac{u_k}{n}(1-f)\right)^\lambda \\ &\geq 1 - (1 - \beta(1-f))^\lambda \geq 1 - e^{-\beta\lambda(1-f)}. \end{aligned}$$

The length of the path established at time I_0 is at least the number of successes in the binomial distribution $B(\tau_0, p)$. Due to the Chernoff bound [2, 12], this length is at least $(1-\gamma)\tau_0 p$ with probability at least $1 - \exp(-\frac{\gamma^2}{2}\tau_0 p)$ for any $\gamma \in (0,1)$. However, $\tau_0 p \geq \alpha \log n \cdot (1 - \exp(-\beta\lambda(1-f))$ and since $\lambda(1-f) \geq 4\ln 2$, we have $\tau_0 p \geq \alpha \log n \cdot (1 - 2^{-4\beta})$, which in turn implies that $\tau_0 p \geq \frac{3}{4}\alpha \log n$, for any $1/2 \leq \beta < 1$. Hence, the length of the traceability path is at least $(1-\gamma)\frac{3}{4}\alpha \log n$ with probability at least $1 - \exp(-\frac{\gamma^2}{2}\frac{3}{4}\alpha \log n) \geq 1 - n^{-\frac{3\alpha\gamma^2}{8}} \geq 1 - n^{-d}$, for any $d \geq \frac{3\alpha\gamma^2}{8} > 1$. ∎

Let E_1 be the event that the path length established during I_0 is at least $\mu_0 = (1-\gamma)\frac{3}{4}\alpha \log n$. Let a *round* of S_1 be the time period until the protocol issues a special attack in order to restart. Consider the sequence of rounds, R, during which $u_k \geq \frac{n}{2}$. Then, for these rounds

$$q = \left(1 - \frac{u_k}{n}(1-f)\right)^\lambda \leq \left(1 - \frac{1-f}{2}\right)^\lambda \leq e^{-\frac{\lambda(1-f)}{2}} \leq 2^{-2} = \frac{1}{4}$$

since $\lambda(1-f) \geq 4\ln 2$. Consequently, $p \geq \frac{3}{4}$.

Let us now condition on the event E_1. For any path length $\mu \geq \mu_0$ consider the path of length μ' derived by "shortcutting" all nodes whose red attacks have all been tried, i.e., by connecting the predecessor of each such node in the path with its successor and thus removing all such nodes from the path. Roughly speaking, μ' represents the number of nodes which can be used for further extending the traceability path. The following lemma shows that $\mu' = \Omega(\mu)$ with high probability.

Lemma 2 *Conditioned on event E_1, the length μ' for the sequence of rounds R is at least $(1-\rho)\frac{3}{4}\mu$, for any $\rho \in (0,1)$, with probability at least $1-n^{-d'}$, $d' > 1$.*

Proof. The length μ' is at least the number of successes in the binomial distribution $B(\mu, \frac{3}{4})$ which, due to the Chernoff bound, is at least $(1-\rho)\frac{3}{4}\mu$ with probability at least $1-\exp\left(-\frac{\rho^2}{2}\frac{3}{4}\mu\right)$, for any $\rho \in (0,1)$. By Lemma 1, this probability is at least $1-\exp\left(-\frac{\rho^2}{2}\frac{3}{4}(1-\gamma)\frac{3}{4}\alpha \log n\right) \geq 1-n^{-\frac{9\alpha\rho^2(1-\gamma)}{32}} \geq 1-n^{-\frac{3\rho^2(1-\gamma)}{4\gamma^2}}$ which is at least $1-n^{-d'}$, $d' > 1$, if we choose γ to satisfy $3\rho^2(1-\gamma) > 4\gamma^2$. ∎

Let E_2 be the event that, conditioned on E_1, $\mu' \geq (1-\rho)\frac{3}{4}\mu$. Then clearly, for the sequence R of rounds and conditioned on both events E_1 and E_2, the process of the path length dominates stochastically a random walk W of length w in the integers $[0,\infty)$, where initially $w=0$ and the probability of a unit move to the right equals $\frac{3}{4}$ while the probability of a unit move to the left equals $\frac{1}{4}$. The next lemma characterizes the above process.

Lemma 3 *Given events E_1, E_2 and for the sequence of rounds R, the number of times S_1 will return to the origin of the established traceability path is bounded above by a constant c, with probability at least $1-e^{-an}$, for some constant $a > 0$.*

Proof. Clearly, it suffices to study the random walk W which is stochastically dominated by the path length process. Let X_i be a random variable taking the value 1 with probability $p = 3/4$ and -1 with probability $q = 1/4$. Then, the length w_k of the random walk W at time k is given by $w_k = w_0 + \sum_{i=1}^{k} X_i = \sum_{i=1}^{k} X_i$ (since $w_0 = 0$). Consequently, the expected length of W at time $n/2$ is

$$E[w_{n/2}] = E\left[\sum_{i=1}^{n/2} X_i\right] = \sum_{i=1}^{n/2}\left(\frac{3}{4}-\frac{1}{4}\right) = \frac{n}{4}.$$

From the Chernoff bound, $\Pr[w_{n/2} \geq (1-\beta)\frac{n}{4}] \geq 1-\exp(-\beta^2 n/8)$, for any $\beta \in (0,1)$.

Let q_z denote the probability that the random walk W starting at time z returns to the origin, and let r denote the number of returns to the origin. Then, from [5, p. 347], $q_z = (q/p)^z = (1/3)^z$, and consequently

$$E[r] = \sum_{z=0}^{\infty} q_z = \sum_{z=0}^{\infty} (1/3)^z = \frac{3}{2}.$$

Now, the probability that W achieves length at least $(1-\beta)\frac{n}{4}$ at time step $n/2$ and then does not return to the origin is $(1-\exp(-\beta^2 n/8)) \cdot (1-(1/3)^{n/2}) = 1-\exp(-\beta^2 n/8) - \exp(-\frac{n}{2}\ln 3) + \exp(-\frac{n}{8}(\beta^2 + 4\ln 3)) \geq 1 - 2\exp(-\beta^2 n/8) = 1-\exp(-an)$, where $a = \beta^2 \ln 2/8$. ∎

We are now ready to give the second main result of this section.

Theorem 2 *The expected number of special attacks performed by S_1 for restarting purposes is bounded above by a constant, provided that $\lambda(1-f) \geq 4\ln 2$ where $\lambda = g/2$.*

Proof. Let s be the number of times S_1 returns to the origin. Clearly, the expected number of special attacks equals $E[s]$. By Lemmata 1–3, it follows that

$$\begin{aligned} E[s] &\leq c\,(1-e^{-an})\Pr[E_1 \cap E_2] + \frac{n}{4}e^{-an}\Pr\left[\overline{E_1 \cap E_2}\right] \\ &< c\,(1-e^{-an})\Pr[E_2|E_1]\Pr[E_1] + 1 \\ &= O(1) \end{aligned}$$

∎

3.3 The attack scheme S_2

The result of Theorem 2 allows the development of another attack scheme S_2 which does not need the special attack to restart the protocol. S_2 simply *simulates* this special attack using a constant number of attacks per node. Protocol S_2 is basically S_1 with the following modifications. It groups the g attacks per node into three equally sized sets, i.e., it now holds that $g = 3\lambda$. The third batch of λ attacks is kept for restarting purposes. To implement the protocol, the node having the token maintains a pointer to its (unique) predecessor and pointers to its successor nodes in the attack. When the path shrinks to a single node, then that node notifies all nodes of the maximum traceability path seen in the past to try to use their third batch of attacks in order to restart the protocol. To find the maximum traceability path, the origin x_0 of the attack can apply a wave algorithm [20, Ch. 6] requesting each node to report its distance in G from x_0. When this information is collected, x_0 requests from the node with maximum distance to initiate the restart process. The properties of S_2 are summarized in the next theorem.

Theorem 3 *Let $\lambda = g/3$, $\lambda(1-f) \geq 4\ln 2$ and $r = \lambda(1-f)/n$. Then: (i) For attack scheme S_2, there exists a time $t_0 = n - e^{-rn}/r + o(n)$ such that if $T \geq t_0$, then $n_{S_2}(T) \geq n(1 - 3\ln 2/g(1-f))$ and $\ell_{S_2}(T) \geq n(1 - 6\ln 2/g(1-f))$ with high probability. (ii) The total failure probability to restart S_2 is $o(1/n)$.*

Proof. Notice that the stochastic process by which u_k changes in S_2 is the same as that of protocol S_1, since at each attempt to extend from a node, a batch of λ attacks is always used. The special attacks which may be required by S_1 for restarting purposes can be simulated, with high probability, by an extra, fixed number of λ attacks per node, as Lemmata 1–3 and Theorem 2 show. Hence, part (i) of the theorem follows by Theorem 1 and the fact that $g = 3\lambda$.

For the proof of part (ii), observe that the total number of attacks that can be used for restarting purposes is $\lambda \cdot L$, where L is the length of the maximum traceability path. By Lemma 1, this is with high probability at least $\delta \log n$ for any constant δ above a certain value. The probability of failure to restart the protocol is then just the probability of these extra attacks all failing to hit U_k. But this is at most

$$\left(1 - \frac{u_k}{n}(1-f)\right)^{\delta \log n} \leq e^{-\frac{u_k}{n}(1-f)\delta \log n}.$$

Since at time k we have that $u_k \geq n - k$, then for $k = t_0$, we get that

$$u_k \geq \frac{e^{-rn}}{r} - o(n) \geq \frac{n}{2\lambda(1-f)e^{\lambda(1-f)}}.$$

Hence, the failure probability to restart the protocol is less than

$$e^{-(\delta \log n)/(2\lambda e^{\lambda(1-f)})} \leq n^{-\delta/(2\lambda e^{\lambda(1-f)} \ln 2)}$$

which is $o\left(\frac{1}{n^2}\right)$ by choosing $\delta > 4\lambda e^{\lambda(1-f)} \ln 2$. Consequently, the total failure probability to restart the protocol is $o\left(\frac{1}{n}\right)$. ∎

The condition $\lambda(1-f) \geq 4\ln 2$ imposes that $g(1-f) \geq 12\ln 2$. Hence, Theorem 3(i) implies that for any $g > g_o = 12\ln 2/(1-f)$ the traceability factor achieved by S_2 is greater than $n/2$, with high probability.

3.4 Three new variants – the tree attack schemes

Instead of keeping the third batch of λ attacks just for restarting S_1 – and thus getting S_2 – we could use them during the path extension process in a more sophisticated way.

Let us first look more carefully at the execution of protocol S_2. S_2 constructs incrementally the graph G whose vertex set is V and whose edges correspond to successful attacks between nodes. The graph G consists of isolated vertices (corresponding to the sleeping nodes), and of a single connected component Δ whose vertices are the awake nodes and whose edges are the edges of G. Since the edges in G represent successful attacks, i.e., attacks from an active to a sleeping node, Δ is clearly a tree. We designate x_0 as the root of Δ.

The new variants are based on the idea that instead of maintaining in S_2 only the maximum traceability path constructed, we maintain the whole tree Δ. Then, consider various orders of the nodes in Δ for path extension using their third batch of attacks. The different orders will specify the different ways the intruder can use to organize his attacks. Clearly, the maximum traceability path constructed by S_2 is the path of maximum depth in Δ. Under this perspective, S_2 can be viewed as simply considering, for path extension, only the nodes of this maximum depth path. If extensions from those nodes fail, then S_2 stops. Hence, it is natural to try to exploit possible expansions from all nodes in Δ, because in such a case the attack has a bigger front to expand compared to that of a single path.

Our three new variants start (like S_2) by emulating S_1 up to the point where all nodes in the current path have exhausted their green and red attacks (the path has shrunk to a single node) and by constructing on the fly the tree Δ. The difference among them and with S_2 lies in the order the nodes of Δ are considered for path extension.

Our first variant, called `rDFS1`, performs a kind of "reverse" DFS on Δ starting with a path of maximum depth and tries to extend the path from its leaf z using the third batch of attacks. If an attack is successful, that is, the returned node v belongs to U_k, then the protocol proceeds as in S_1: the path is extended, v is equipped with g attacks and takes the token. If all attacks fail, then `rDFS1` backtracks to the parent u of z and repeats the process (i.e., tries to extend the path from u). If during this process all nodes in the path (except for x_0) exhaust their third batch of attacks, then the protocol is restarted from the next largest path in Δ. The whole process is repeated until either $U_k = \emptyset$ or all nodes have exhausted their attacks. At any time during the execution of the protocol the tree Δ is maintained.

Note that as protocol `rDFS1` moves backwards on a path from a leaf z to x_0 and fails to extend the path from an internal node v, the nodes in the subtree rooted at v will be considered at some later point of time by the protocol depending on the overall length of the path to which they belong. An alternative order to consider these nodes is provided by our second variant, called `rDFS2`. It is similar to `rDFS1` except that when the protocol moves backwards to an internal node v, it doesn't try to extend the path from v but instead moves to the leaf z' belonging to a maximum depth path in the subtree of Δ rooted at v and tries to extend the path from z'.

Our third variant, called `rBFS`, is similar in spirit with the above two, but it employs a different strategy in the way the nodes of Δ are considered for extending the path. Protocol `rBFS` performs a kind of "reverse" BFS on Δ starting with the nodes of maximum depth d. It tries to extend the path by considering all such nodes one-by-one (each node uses its third batch of attacks). If an attack is successful, then the protocol proceeds as in S_1: the path is extended, d is increased by one, the new node returned is equipped with g attacks and takes the token. If all attacks in all nodes at level d fail, then `rBFS` repeats the process with all nodes at the tree level $d-1$. The whole process is repeated until either $U_k = \emptyset$ or all nodes in Δ have exhausted their attacks.

In the following, we shall refer to `rDFS1`, `rDFS2`, and `rBFS` as the *tree attack schemes* or *tree protocols*. Notice that all protocols can be easily implemented in a distributed way. To implicitly maintain the tree Δ, every node keeps a pointer to its parent and pointers to its children. If the node z_k having the token succeeds, then it passes it to the next node z_{k+1} and keeps a pointer to z_{k+1} (new child). Otherwise, it passes the token to its parent, say u. Depending on the protocol which is currently executed, u either starts issuing attacks, or simply forwards the token to the appropriate node which can be found by applying a wave algorithm [20, Ch. 6] (as in the case of S_2).

Despite their apparent similarity, we investigate in the remainder of the paper all tree attack schemes, since it turns out that they either differ in terms of analysis or in terms of experimental behaviour.

3.5 Structural analysis of the tree attack schemes

Notice that the stochastic process by which u_k changes in a tree attack scheme S is again the same as that of protocol S_1, since at each attempt to extend from a node, a batch of λ attacks is always used. Hence, $n_S(k) = n - u_k = |\Delta_k|$, where Δ_k is Δ at time k. Also, $\ell_S(k) = \ell_k$, where ℓ_k is the length of the traceability path. Analogously to the analysis of S_1 (Section 3.2), we would like to establish a lower bound on ℓ_k as a function of n, k and u_k.

Let Y_k be the subset of red nodes that belong to Π_k and are reused to extend the path via their third batch of attacks in the tree attack schemes (i.e., each node in Y_k has already used all of its red attacks), and let $y_k = |Y_k|$. Arguing as before, we get $V - U_k - \Pi_k \subseteq R_k - Y_k$ which implies that $\ell_k \geq n - u_k - r_k - 1 + y_k$. Combining this with Eq. (1), we obtain

$$\ell_k \geq 2(n - u_k) - (k+2) + y_k \tag{5}$$

Comparing the above equation with (2), we see that the tree attack schemes add (at least) the quantity y_k to the path length achieved. Obtaining a precise estimation for y_k is very hard, but fortunately we can proceed in an alternative way. Recall that one of our interests is to study the relationship between ℓ_k and $|\Delta_k|$. Before diving into an analytic study, we performed a series of experiments to obtain it experimentally. This would give us an indication and, as we shall see later, provide us with a basis for some reasonable assumptions required by the analysis.

Our experiments clearly indicated that there is a constant $1/2 < a < 1$ (actually $a \geq 0.68$) such that $\ell_k \geq a|\Delta_k|$ (see Figure 1), provided that g satisfies $g(1-f) \geq 12\ln 2$ as required by the analysis[2]. This experimental result can be verified analytically in the case of `rDFS1` as we show next.

We first consider the situation where `rDFS1` proceeds successfully (i.e., emulates S_1) without using any of the third batch of attacks. Let a *green node* be an awake node whose green set of attacks has not been exhausted. By working in a way identical to that of the proof of Lemma 1, we can establish the following.

Lemma 4 *Starting from a green node and emulating S_1, attack scheme* `rDFS1` *will extend the traceability path to an additional length $\ell^* = \Theta(\log n)$, without using any of the third batch of attacks, with probability at least $1 - n^{-c}$ (for some constant $c \geq 2$), provided that $\lambda(1-f) \geq 4\ln 2$ where $\lambda = g/3$.*

Let us now consider what happens when the protocol is forced to use the third batch of attacks. Then, the current traceability path will "backtrack" (because of failures in the third batch) until a node extends again to another direction. The backtracking probability is

$$q(k) = \left(1 - \frac{u_k}{n}(1-f)\right)^{\lambda}$$

which is bounded from above and below by constants $q' \leq q(k) \leq q$ (independent of n) as long as $u_k = \Theta(n)$. But in such a case the backtracking process is dominated by a geometric process of a

[2] We have observed, however, that $a > 1/2$ even for smaller values of g.

constant backtracking probability q for each backtracking step. The probability that the additional path ℓ^* (gained by the sequence of extensions) will then drop to $\ell^*/2$ is at most $q^{\ell^*/2} \leq n^{-c_1}$, for some suitably chosen constant $c_1 \geq 2$.

We shall use the above fact along with Lemma 4 to establish the desired relationship between ℓ_k and $|\Delta_k|$. Let us condition on all times during which $u_k = \Theta(n)$ and let us consider the i-th part of the process of extending the path to an additional length ℓ^* and then backtrack. Call this i-th part π_i. Let A_i be the event "the extension to an additional length ℓ^* gives $\ell^* = \Theta(\log n)$ and the backtracking reduces it to at most $\ell^*/2$". Then, we have that

$$\Pr[A_i] = (1 - n^{-c})(1 - n^{-c_1}) \geq 1 - 2n^{-c_2}$$

where $c_2 \geq 2$. Now, consider $n/2\log n$ π_i's. The probability that at least one A_i does not hold in this sequence is

$$\Pr[\exists \overline{A_i}] \leq \sum_i \Pr[\overline{A_i}] \leq \left(\frac{n}{2\log n}\right)(2n^{-c_2}) < \frac{1}{n}.$$

Hence, all events A_i hold simultaneously with probability at least $1 - 1/n$, i.e., the backtracking in each π_i at most halves the extension of the path achieved within π_i. Clearly, the cumulative path extensions in all π_i's equals $|\Delta_k|$. Since the cumulative backtrackings in all π_i's reduce $|\Delta_k|$ by at most a half, we conclude that $\ell_k \geq |\Delta_k|/2$. The preceding discussion, along with the fact that the stochastic process by which u_k changes is the same as that of protocol S_1, establishes the following.

Theorem 4 *Let $\lambda = g/3$ and $\lambda(1-f) \geq 4\ln 2$. Then, with high probability, there exists a time step t_0 for attack scheme* `rDFS1` *such that: (i) $\forall 0 \leq k \leq t_0$, $\ell_{\texttt{rDFS1}}(k)/n_{\texttt{rDFS1}}(k) \geq 1/2$; (ii) $\forall T \geq t_0$, $n_{\texttt{rDFS1}}(T) = \Theta(n) = \ell_{\texttt{rDFS1}}(T)$.*

Unfortunately, the above analysis does not carry over to attack schemes `rDFS2` and `rBFS`. However, our experimental results (cf. Figure 1) indicate that they perform at least as good as `rDFS1`. Hence, our experimental evidence and the above theoretical analysis lead us to the following experimental conclusion (or conjecture for `rDFS2` and `rBFS`) .

Experimental Conclusion (Conjecture) 1 *With high probability, there exists a time step t_0 for any tree attack scheme S such that: (i) there is a constant a, $1/2 \leq a < 1$, such that $\forall 0 \leq k \leq t_0$, $\ell_S(k)/n_S(k) \geq a$; (ii) $\forall T \geq t_0$, $n_S(T) = \Theta(n) = \ell_S(T)$.*

We shall use the above conclusion as a (reasonable) hypothesis in order to investigate analytically the relationship between the traceability factor achieved by the tree protocols and the one achieved by S_1 or S_2.

Recall that $n_S(k) = n - u_k$ and that the stochastic process by which u_k changes is the same as that of protocol S_1. We will now condition on the event "$\forall k, \ell_k \geq a(n - u_k)$" and provide evidence that each of the tree protocols achieves slightly larger traceability factors than that of S_1 and S_2. This is also verified by our experiments.

Theorem 5 *The tree attack schemes achieve a traceability factor which is at least as large as the one achieved by attack schemes S_1 and S_2, provided that Experimental Conclusion 1 is true.*

Proof. Let $\tilde{u}_k$ be any upper bound on u_k holding with high probability. We only have to compare the progress of two expressions, namely $\mathcal{E}_1(k) = 2(n - \tilde{u}_k) - (k+2)$ for S_1, S_2, and $\mathcal{E}_2(k) = a(n - \tilde{u}_k)$ for the tree protocols.

Let $\lambda(1-f) \geq 4\ln 2$, $r = \frac{\lambda}{n}(1-f)$, and $j = \lceil n\ln 2/\lambda(1-f)\rceil$. Then, we have from Theorems 1, 3, 4, and Experimental Conclusion 1 that the first time step k at which the path length becomes

$\Theta(n)$ is, with high probability, $n - e^{-rn}/r + o(n)$. For this value of k, we have that $\tilde{u}_k \leq j$ and the above expressions become

$$\mathcal{E}_1(k) = 2\left(n - n\ln 2/\lambda(1-f)\right) - \left(n - \frac{e^{-rn}}{r} + o(n)\right)$$

and

$$\mathcal{E}_2(k) = a\left(n - n\ln 2/\lambda(1-f)\right).$$

Since $\lambda(1-f) \geq 4\ln 2$, it is a matter of simple calculations to check that $\mathcal{E}_2(k) - \mathcal{E}_1(k) > 0$ for any $a \geq 1/2$. ∎

Notice that this is only a lower bound comparison, since both our analysis for S_1 and its extension to the tree attack schemes are not giving a precise value for the traceability factor, but only lower bounds.

4 Experimental Results

As it is customary with such protocols, our implementations were done on a simulation environment. In the implementations we tried to exploit the full strength of the protocols, in order to investigate their practicality. The (pure) protocol statements discard the remains of batches in attacks which succeeded to extend the traceability path, in the case where a part of the batch was processed. In our implementation we also used these remains. The implementations of all protocols were written in `C++` and were run on a SUN Enterprise 450 running Solaris 2.7 and having 1.2GB of main memory.

We performed an extensive set of experiments using different types and sizes of inputs. We considered networks with a number of vertices n as small as 50 and as large as 10000. For each such network we considered several values for the critical parameters of the problem, namely f (failure probability) and g (number of available attacks per node), ranging from 0.1 up to 0.95 for f, and from 2 up to 200 for g. We report results with representative values of n (5000), and of the critical parameters f and g (the conducted experiments with other values of n, f and g reported similar results). We consider two values for f, namely 0.75 (network of above-average security level), and 0.95 (highly secure network), and two values for g, namely 10 and 50. We generated a large collection of data sets, each consisting of 50 samples. Each data set corresponded to a fixed value for the network parameters n, f, and g. All protocols were run on each sample and the reported performance measures are averages over the 50 samples. The conducted experiments confirmed our theoretical results and exhibited the practicality of the protocols.

We measured the length L of the maximum traceability path achieved by some protocol either as a function of g given a fixed value for f (Figure 2), or as a function of f given some fixed value of g (Figure 3). The `Original` curve refers to the implementation of protocol S_2. The `Theory` curve shows the lower bound expected for the maximum traceability factor from Theorem 3. Note that this lower bound holds only when $g \geq g_o = 12\ln 2/(1-f)$. Consequently, the zero values in the `Theory` curve represent the trivial lower bound that holds in the cases where $g < g_o$, since in such cases the analysis does not provide any non-trivial lower bound. It is interesting to observe that for small values of g, all protocols achieve a much better traceability factor than `Theory`, even when f gets larger. When f is relatively small (e.g., $f \leq 0.75$), all protocols achieve almost the same maximum traceability factor, mainly because the number of restarts is minuscule. When f gets larger, however, `rBFS` usually achieves a better path length and is the most robust (having much less variance than the other protocols).

We observe no big differences w.r.t. the maximum traceability factor achieved by the protocols. This can be explained by the fact that in all experiments we performed the resulting tree Δ was almost always deep and lean. This implies that there are neither big differences among the various orders in which the protocols consider the nodes of Δ, nor big differences between `Original` and the tree protocols, and this carries over to the length of the maximum traceability path. For very large values of f, e.g., $f = 0.95$, we observed a big variance below some threshold of g (smaller than g_o). While it seems difficult to explain this theoretically, we suspect that it is due to the following. If the first few nodes fail to issue their attacks (which is highly probable in this case), then the length of the path will be very small, and as a consequence it is highly improbable that the protocol will even be able to restart. On the other hand, once a protocol manages to restart, the achieved path length is much better than the theoretically expected one – even if a relatively larger number of restarts is required (see Figure 4 and the discussion below regarding restarts). We also observe that in such cases the protocols need much less than g_o attacks to achieve a reasonable path length, thus demonstrating their good practical behaviour.

Regarding the number of restarts, our experiments showed that all protocols perform *zero* restarts for any value of n we considered ($n \in [50, 10000]$) when $g > 45$ and relatively small f (i.e., $f \leq 0.75$); see Figure 4 (left diagram). The number of restarts may vary almost linearly with n as g approaches g_o (see Figure 4), but drops rather quickly as soon as it passes g_o. The big variance observed in the number of restarts (cf. right diagram of Figure 4) for large values of f and when g is below g_o is due to the length of the maximum traceability path achieved by the protocol prior to the first restart. If this length is small, then the protocols perform very few restarts, after which they stop (since all available attacks have been exhausted). Otherwise, the protocols, due to the high probability failure, perform many restarts in order to extend the path. This explains why e.g., `Original` has much less restarts than any of the tree protocols.

5 Conclusions and Further Research

We investigated the analytic and experimental behaviour of four protocols for the attack propagation problem in networks under a new model for intrusion propagation introduced here. We have shown that a limited in power intruder can spread his attack in a large part of a network, given sufficient propagation time, and that a detection mechanism has to trace a number of nodes proportional to the network size in order to find the origin of the intrusion. An obvious direction for further research is to extend our analysis so that the exact dependence of the spread and traceability factors on time is estimated at all time instances (and not only for times larger than a given quantity).

Our model is a first attempt towards more realistic modeling of intrusion propagation and it would be interesting to investigate further extensions of it. Perhaps one limitation of our model is the way it represents the cases of failure of an attack. We model this by a single fixed parameter f. However, the actual resistance to attacks may vary from system to system and also it may vary with time; for example, the longer the attack spreads, the more difficult it may become to break other systems, since they may dynamically increase their defenses. Furthermore, the defense ability of a system may not depend on a single parameter only, but on several parameters (e.g., perimetric security devices, available countermeasure software, experience of the security manager, etc). Therefore, one interesting research direction is to strengthen the analysis by considering the parameter f to depend on system and time.

Another extension of our model could be the generalization of the notion of traceability to involve not only the distance from the origin of the intrusion, but also the ability of the defense mechanism to follow backward links. The latter may not be easy in cases where these links are encrypted by the intruder. A somehow related, but complementary, work for the ability to trace back IP packets to an anonymous origin has very recently appeared in [1]. In that work, a sophisticated way is proposed to encode the distance from the origin into the packet header bits.

A natural choice for the intrusion propagation is to select the target nodes (systems) to attack uniformly at random. Such a process has also been followed in other lines of research, for example in web graph models [3], which consider dynamic graphs created by the uniform choice of a new vertex at each time step along with some edges. The models in [3] lead to graphs similar to the graphs created by our attack schemes (where edges representing non-successful attacks are also included in G). It would be interesting to investigate attack propagation schemes based on non-uniform selection of target nodes.

Finally, our model may also find applications in the context of dynamic games in network computing. The construction of the graph G can be viewed as a game in which edges are added to G under certain conditions (those described in Section 2) and where one wishes to optimize various parameters in G (e.g., size of connected components, distance between specific nodes, etc). Such dynamic games have been introduced in the context of network computing by Peleg [16, 17] and continued by many others [6, 7]. These papers investigate optimal values for the sizes of the so-called dynamic monopolies (dynamos) which are patterns of initial faults in a network (that spread using some majority rule based on the faults of neighboring nodes) and whose occurrence could lead the entire system to a faulty behaviour (catastrophe).

Acknowledgments

We would like to thank the anonymous referees for their valuable comments and criticism which improved the presentation of the paper.

References

[1] M. Adler, "Tradeoffs in Probabilistic Packet Marking for IP Traceback", in *Proc. 34th ACM Symp. on Theory of Computing* – STOC 2002 (ACM Press 2002), pp.407-418.

[2] H. Chernoff, "A Measure for Asymptotic Efficiency for Test of a Hypothesis Based on the Sum of Observations", *Annals of Math. Stat.*, 23 (1952), pp.493-509.

[3] C. Cooper and A. Frieze, "Crawling on Web Graphs", in *Proc. 34th ACM Symp. on Theory of Computing* – STOC 2002 (ACM Press 2002), pp.419-427.

[4] D. Denning, *Information Warfare and Security*, Addison-Wesley, 1999.

[5] W. Feller, *An Introduction to Probability Theory and its Applications*, Vol. I, John Wiley, New York, 1968.

[6] P. Flocchini, F. Geurts, and N. Santoro, "Optimal Irreversible Dynamos in Chordal Rings", in *Proc. 25th Workshop on Graph-Theoretic Concepts in Computer Science* – WG'99, *Lecture Notes in Computer Science* Vol. 1665 (Springer-Verlag, 1999), pp.202-214.

[7] P. Flocchini, E. Lodi, F. Lucio, L. Pagli, and N. Santoro, "Irreversible Dynamos in Tori", in *Proc. 4th Euro-Par (Parallel Processing) Conference* – Euro-Par'98, *Lecture Notes in Computer Science* Vol. 1470 (Springer-Verlag, 1998), pp. 554-562.

[8] J. Howard, "An Analysis of Security Incidents on the Internet 1989-1995", Carnegie Mellon Univ., CERT/CC, 1997, `www.cert.org/research/JHThesis/Start.html`.

[9] J. Kephart and S. White, "Directed-Graph Epidemiological Models of Computer Viruses", in *Proc. IEEE Symp. on Security and Privacy*, 1991, pp.343-359. Also IBM Research Report: `www.research.ibm.com/antivirus/SciPapers/Kephart/VIRIEEE/virieee.gopher.html`.

[10] J. Kephart and S. White, “Measuring and Modeling Computer Virus Prevalence”, in *Proc. IEEE Symp. on Security and Privacy*, 1993, pp.2-14. Also IBM Research Report: `www.research.ibm.com/antivirus/SciPapers/Kephart/PREV/prevalence.gopher.html`.

[11] T.M. Liggett, *Stochastic Interacting Systems: Contact, Voter and Exclusion Processes*, Springer-Verlag, 1999.

[12] R. Motwani and P. Raghavan, *Randomized Algorithms*, Cambridge University Press, 1995.

[13] S. Nikoletseas and P. Spirakis, “Efficient Communication Establishment in Adverse Communication Environments”, in *Proc. Workshop on Approximation and Randomized Algorithms in Communication Networks* – ARACNE 2000 (ICALP 2000 Satellite Workshop), Ed. J. Rolim, et al, Carleton University Press, 2000, pp.215-226.

[14] S. Nikoletseas, G. Prasinos, P. Spirakis, and C. Zaroliagis, “Attack Propagation in Networks”, in *Proc. 13th ACM Symposium on Parallel Algorithms and Architectures* – SPAA 2001 (ACM Press, 2001), pp. 67-76.

[15] R. Ostrovsky and M. Yung, “How to Withstand Mobile Virus Attacks”, in *Proc. 10th ACM Symp. on Principles of Distributed Computing* – PODC’91 (ACM Press, 1991), pp.51-59.

[16] D. Peleg, “Local Majority Voting, Small Coalitions and Controlling Monopolies in graphs: A Review ”, in *Proc. 3rd Colloquium on Structural Information and Communication Complexity*, Carleton University Press, 1996, pp.152-169.

[17] D. Peleg, “Size bounds for Dynamic Monopolies”, in *Discrete Applied Mathematics*, 86 (1998), pp. 263-273.

[18] T. Shimonura and J. Markoff, *Takedown: The Pursuit and Capture of Kevin Mitnick, America’s most Wanted Computer Outlaw*, Hyperion, NY 1996.

[19] D. Safford, D. Shales, and D. Hess, “The TAMU Security Package: An Ongoing Response to Internet Intruders in an Academic Environment”, in *Proc. UNIX Security Symposium IV*, 1993.

[20] G. Tel, *Introduction to Distributed Algorithms*, Cambridge University Press, 2nd edition, 2000.

[21] S. White, “Open Problems in Computer Virus Research”, in *Virus Bulletin Conference*, 1998. IBM Research Report: `www.research.ibm.com/antivirus/SciPapers/White/Problems/Problems.html`.

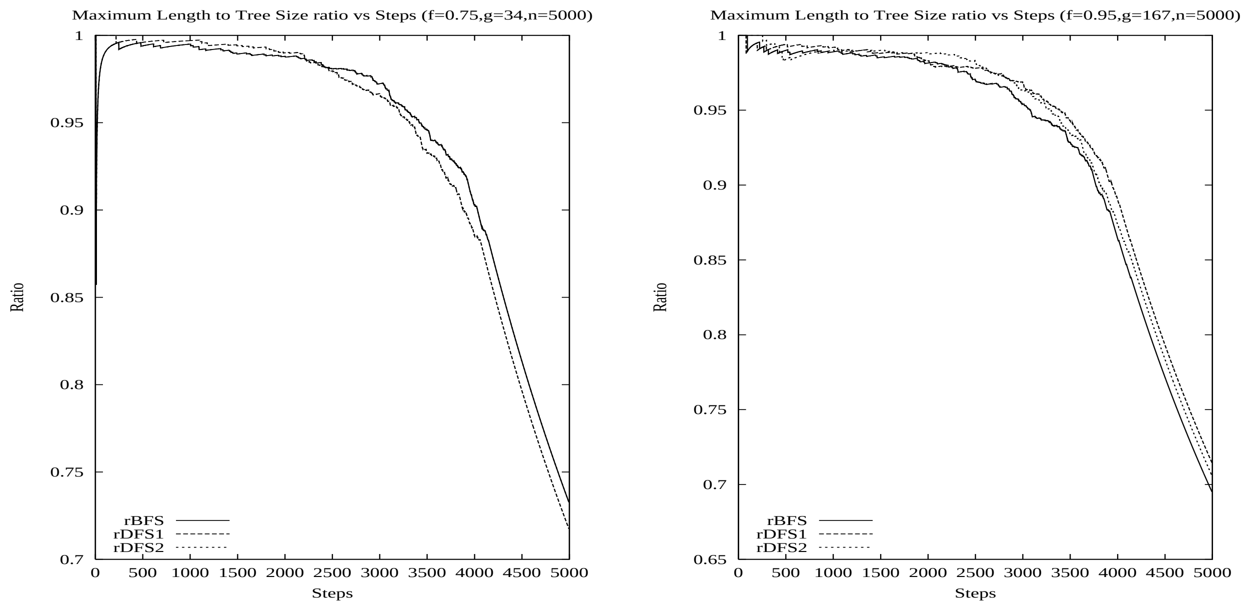

Figure 1: The ratio $\ell_k/|\Delta_k|$ versus the number of steps for which extension of the traceability path is achieved.

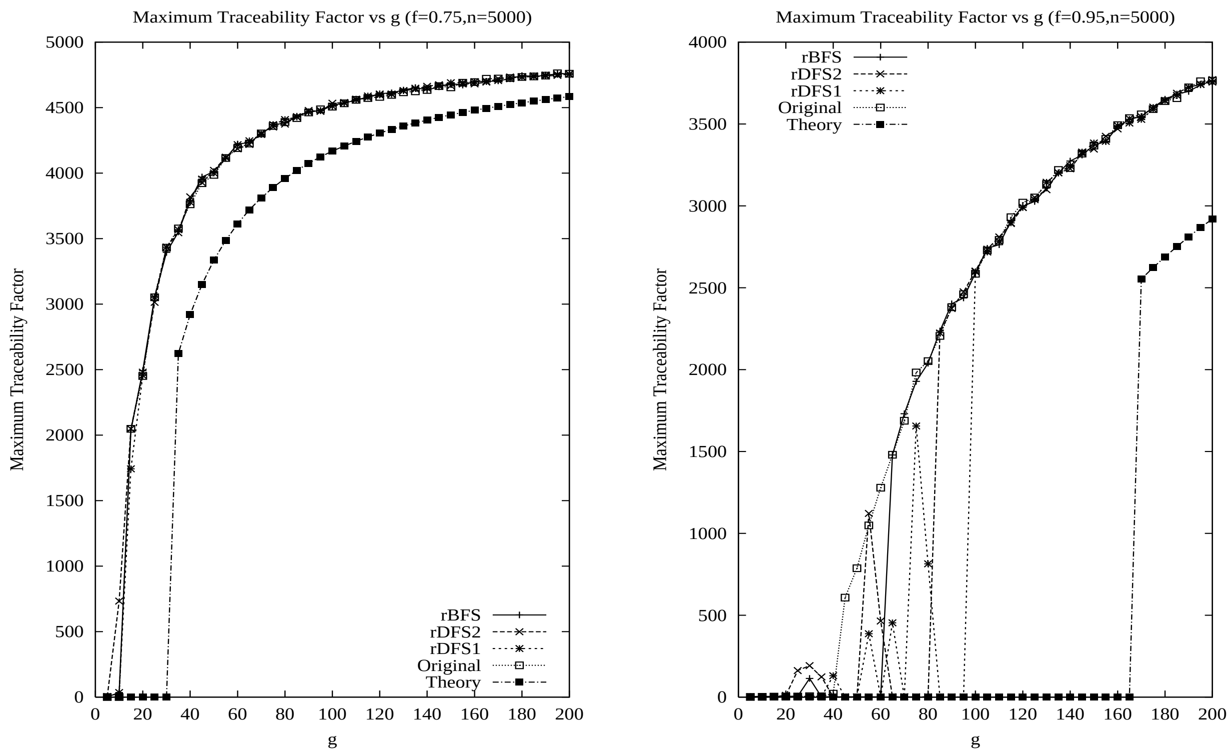

Figure 2: Maximum traceability factor as a function of g for $n = 5000$. The granularity of the g-axis is 5.

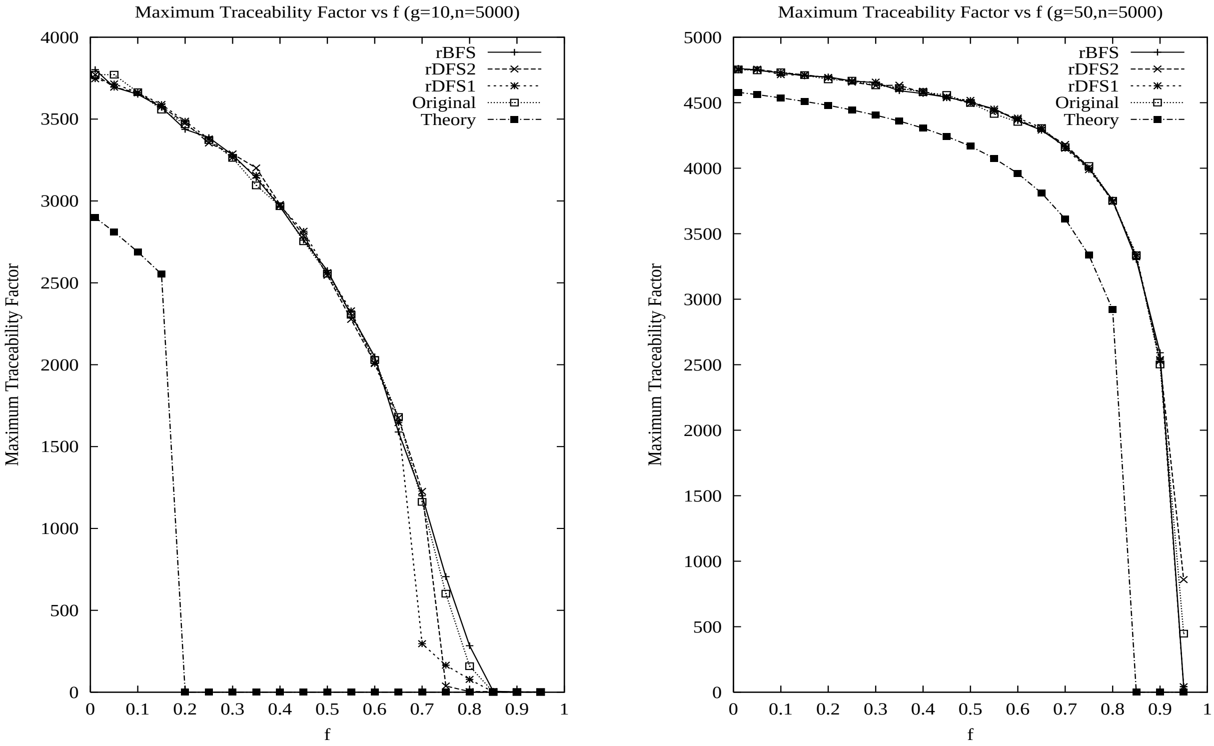

Figure 3: Maximum traceability factor as a function of f. The granularity of the f-axis is 0.05.

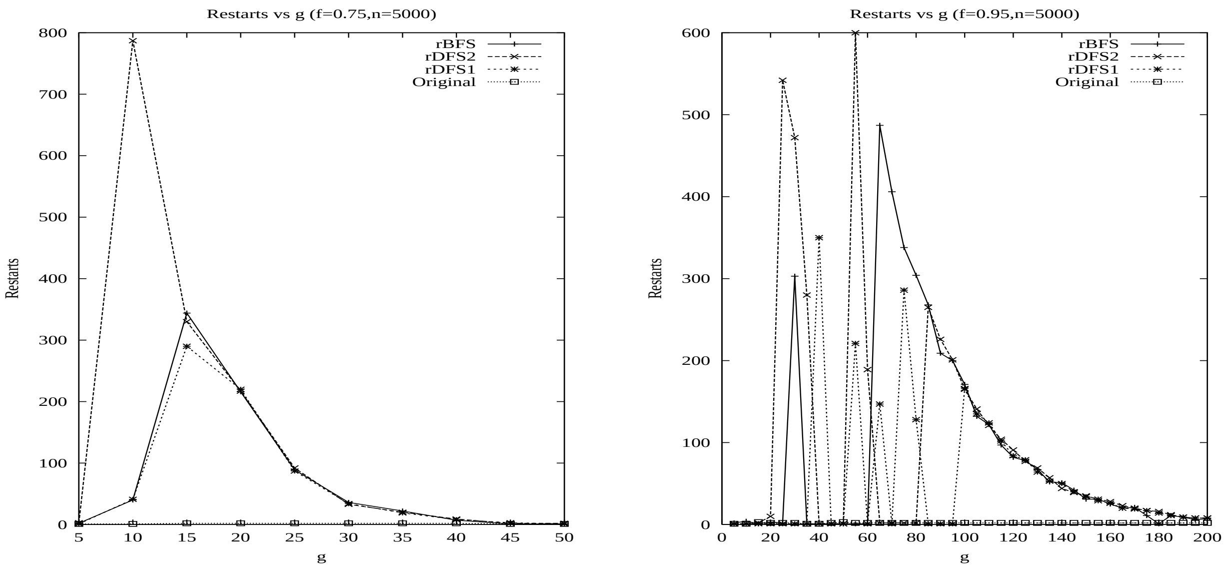

Figure 4: Number of restarts for $f = 0.75$ and $f = 0.95$ ($n = 5000$).