Privacy and Reliability by Dispersive Routing

Haim Zlatokrilov and Hanoch Levy

School of Computer Science Tel-Aviv University, Tel-Aviv, Israel

Abstract — The traditional single (shortest) path routing paradigm leaves sessions vulnerable to a variety of security threats, especially eavesdropping. We propose to overcome this via dispersive routing, conducted over multiple paths. This increases significantly the costs inflicted on an attacker who wishes to eavesdrop or conduct DoS attack on network sessions by hijacking network links (or routers)[1].

Introduction

The traditional single (usually shortest) path routing leaves sessions vulnerable to attacks along the route. Attackers may eavesdrop sessions as well as maliciously drop their fragments (causing denial-of-service (DoS) attack), on nodes or links along the path. The approach proposed in this work is to enhance privacy and reliability by adding additional layer of protection. While encryption is a good defense against attackers that managed to eavesdrop an entire session, the dispersion of session fragments over multiple paths can prevent the attacker from conducting a meaningful eavesdropping or significant malicious dropping[2] in the first place.

Our model is based on the assumption that each link is associated with some adversary hijacking cost. This cost is based on parameters such as physical link properties (e.g. wire or wireless), geographic location, etc. We study the problem of shipping session fragments in a way that will force the attacker to invest at least a predefined minimal effort to conduct a successful attack. We look at the worst-case scenario, assuming the attacker is familiar with the exact dispersion strategy and knows the path taken by each fragment. Comprehensive study of this problem with several extensions, such as finding minimal number of paths and limiting paths' length, will be presented in [4].

Dispersing session fragments over multiple paths can be implemented in a variety of methods such as: IP tunneling, implementation in overlay or MPLS networks, etc. We assume that the Security Traffic Manager (STM) can plan and execute the transmission of session fragments regardless of the underlying machinery. We focus on the context of a single session and neglect bandwidth constraints, assuming that session's bandwidth requirements[3] are very small in comparison to network capacity.

[1] This work was partially supported by Israeli Science Foundation grant 235/03 and by Euro NGI network of excellence.

[2] We will use the term *dropping* both for eavesdropping and malicious dropping attacks.

[3] Dispersion techniques are also known to increase network efficiency as discussed in [3].

Sending session fragments over several paths might cause degradation in QoS due to jitter and out-of-order effects, but in some scenarios can also enhance QoS [1].

A Brief Demonstration of the Problem

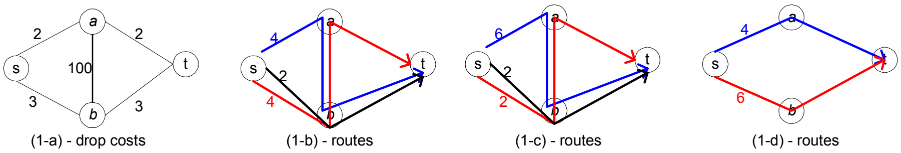

Fig. 1. An example of the STM's problem. (1-a) - network structure and link costs. (1-b) and (1-c) - bad assignment of fragments over the paths (the numbers depict the number of assigned fragments over the paths). (1-d) – a good assignment which protects the session.

Consider the network depicted in Fig.1-a, where the numbers represent link hijacking costs. The STM's goal is to transfer a session of 10 fragments from *s* to *t*, such that the attacker will be forced to invest at least cost of 5 units in order to drop 8 fragments. If the fragments are directed as shown in Fig.1-b or Fig.1-c, the attacker will be able to drop 8 fragments at the cost of 4 units by hijacking links *s*→*a* and *a*→*t* (or link *b*→*t* in Fig.1-c). Using the dispersion strategy depicted in Figure.1-d will keep the session safe.

The Attacker's Problem

INSTANCE: Graph *G*(*V*,*L*), cost c_l for all links $l \in L$, a set p_i, $i = 1..K$, of paths from *s* to *t*, the path taken by each of *N* fragments and a parameter *P*. Let $\gamma_i^j = 1$, denote that fragment *j* is sent over path p_i, otherwise $\gamma_i^j = 0$ (Clearly $\sum_{j=1}^{N}\sum_{i=1}^{K}\gamma_i^j = N$).

QUESTION: Find a set of edges $L' \subseteq L$, such that $\sum_{l \in L'} c_l$ is minimized and the number of unique[4] fragments on L' is greater than or equals to $P \cdot N$.

Lemma 1: *The attacker's problem is NPC even with identical link costs for any* $0 < P \leq 1$. *The attacker's problem has a (1+ln(*$|V|$*))-approximation by a greedy algorithm.*

The STM's problem

INSTANCE: Graph *G*(*V*,*L*), cost c_l for all links $l \in L$, source and destination nodes (*s* and *t* respectively), a parameter $0 < P \leq 1$ and a constant *C*>*0*.

QUESTION: Find a set of paths *U* from *s* to *t* and an assignment of *N* fragments to the paths such that: There exists no set of links *L*', obeying $\sum_{l \in L'} c_l < C$, that captures $P \cdot N$ fragments.

We prove that the STM's problem in general is at least NP-hard. Nonetheless several special cases (covering large set of parameters) of the problem can be solved in a polynomial time.

The STM's problem with either *P*=1 or *P*<1 and identical cost links

We propose simple polynomial time algorithms solving these cases of STM's problem, where the number of session fragments obeys *N*>*C*/c_{min} (c_{min} is the minimal

[4] A fragment is counted only once, even if the attacker drops it several times.

cost link). If no solution exists the algorithms fail. The algorithms are based on finding max-flow and translating costs to number of fragments to be sent on paths. The complexity of these algorithms is identical to that of the max-flow algorithm.

The STM's problem with P<1 and non-identical cost links

We show that the problem is at least NP-hard. We propose a heuristic for the STM's problem possessing the following properties (C' denotes the minimal cost cut):

- If the attacker's budget $C>C'$, the attacker can always prevail and the algorithm stops.
- If $C < P \cdot C'$, a solution that guarantees that the STM prevails is found.
- Otherwise, if $P \cdot C' < C < C'$, it is not clear whether a valid solution exists. Several enhancement heuristics can be used in that case. Validating the solution in that case, which is the attacker's problem, is NPC.

Simulation Results

Fig.2 demonstrates the increase in the attacker's required dropping budget as a function of node degree. This is achieved by running the *STM's Algorithm* and computing the capturing cost (*heuristic's solution* curve). The simulations were conducted on random graphs with random edge costs in the range of 1-5 and P=0.8. The plot depicts the averaged results. The *max single path cost* curve stands for the maximal dropping cost in case of single path routing. The *min cost cut* curve stands for the budget for which the attacker will always prevail. The figure demonstrates that implementing the *STM's algorithm* can dramatically increase the dropping cost in comparison to single path routing, up to the cost $P \cdot C' < C < C'$

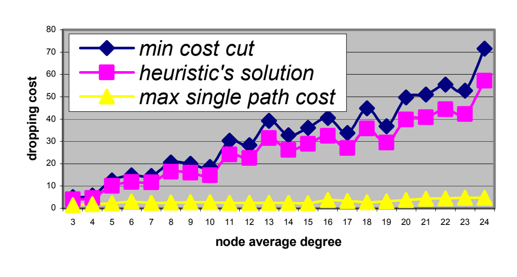

Fig. 2. increase in attacker's dropping budget

References

1. H. Zlatokrilov, H.Levy, "The Effect of Packet Dispersion on Voice Applications in IP Networks", in Proc. of IEEE INFOCOM'04, Hong-Kong.
2. S. Bohacek, J. P. Hespanha, K. Obraczka, J. Lee, C. Lim, "Enhancing Security Via Stochastic Routing", In Proc. of the 11th IEEE ICCCN, May 2002.
3. Patrick P. C. Lee , V. Misra , and D. Rubenstein. "Distributed Algorithms for Secure Multipath Routing", in Proc. of IEEE INFOCOM, March 2005, Miami, FL, USA.
4. H. Zlatokrilov, H. Levy, "Session Security enhancement by traffic dispersion", forthcoming.