Traffic Analysis Attacks in Anonymity Networks: Relationship Anonymity-Overhead Trade-off

Ognjen Vuković
School of Electrical Engineering
KTH, Royal Institute of Technology,
Stockholm, Sweden
Email: vukovic@ee.kth.se

György Dán
School of Electrical Engineering
KTH, Royal Institute of Technology,
Stockholm, Sweden
Email: gyuri@ee.kth.se

Gunnar Karlsson
School of Electrical Engineering
KTH, Royal Institute of Technology,
Stockholm, Sweden
Email: gk@kth.se

***Abstract*—Motivated by applications in industrial communication networks, in this paper we consider the trade-off between relationship anonymity and communication overhead in anonymity networks for traffic analysis attacks. For our study, we use two anonymity networks: MCrowds, a variant of Crowds, that provides unbounded communication delay and Minstrels, that provides bounded communication delay. Our results show that, contrary to intuition, increased overhead does not always improve anonymity.**

I. Introduction

Many communication systems, for example modern industrial networks [1], [2], require high availability between a fixed set of nodes on a pairwise basis. The nodes can be the subsidiaries of an enterprise connected by a virtual private network over the public Internet, or they can be sensors, actuators and operation centres in a wide area industrial control system, e.g., in a supervisory control and data acquisition (SCADA) network. Cryptography may provide authentication, confidentiality and data integrity for the communication, but source and destination addresses could still be visible to an outside attacker who is able to observe one or more network links. The outside attacker may identify traffic patterns: who is communicating with whom, when and how often. Using this information the attacker can infer the importance of the messages, and may perform targeted attacks on the communication between any two nodes. These targeted attacks might be hard to detect and can lead to incorrect system operation.

Mix networks [3] are a way to mitigate outside attacks by providing relationship anonymity, i.e., by making it untraceable who communicates with whom. Nodes in a mix network relay and delay messages such that an outside attacker cannot trace the route of the individual messages through the mix. While relaying renders outside attacks more difficult, it introduces the possibility of inside attacks. Due to the often long life-cycles of industrial systems software corruption is a threat, and the complexity of the code-base makes corruption hard to detect. Corrupted nodes that are part of the mix network can perform inside attacks to determine the sender-receiver pair for messages that are relayed through them. Anonymity networks can be used to provide relationship anonymity against inside attackers (e.g., [5]) by hiding the sender or the receiver from the relay nodes. Nevertheless, the relationship anonymity provided by mix networks and anonymity networks comes at the price of delay and communication overhead. Excessive delays can negatively impact the system performance, while overhead leads to high resource requirements, so that in practice both have to be kept low.

Intuition says that increased overhead (and delay) should result in increased anonymity in anonymity networks. In this work we show that this is not always the case.We consider an attacker whose goal is to perform a traffic analysis attack in order to determine the communication patterns between a set of communicating nodes, i.e., the traffic matrix. We consider two methods for traffic analysis: a *Bayesian inference* and a *Maximum posteriori* method. According to the *Bayesian inference* method the attacker considers all pairs of nodes as a possible sender-receiver pair for an intercepted message. According to the *maximum posteriori method* the attacker only considers the most likely pairs of nodes as a possible sender-receiver pair for an intercepted message. We use two anonymity networks for our study. First, MCrowds, described in this paper, which is an anonymity network very similar to Crowds [5]. MCrowds hides the sender by introducing unbounded message delivery delay, and hides the receiver among a small subset of anonymity network users. Second, Minstrels, proposed in [6], which provides bounded message delivery delay by limiting the maximum number of visited nodes for each message, and hides the sender and the receiver among all anonymity network users.

II. System Model and Metrics

We consider an anonymity network with N nodes. The nodes act as sources, destinations and as relay nodes for each others' messages. The underlying communication network is a complete graph. The *inside attacker* is in control of C nodes, and can observe the messages traversing those nodes and the protocol specific information contained in the messages. The attacker's goal is to identify the source and the destination of the messages that it observes. To achieve its goal, the attacker calculates for every pair of nodes $(a,b) : a \in N, b \in N$ in the system the probability that it is the sender-receiver pair (s,r). For the calculation the attacker assumes an a-priori traffic matrix, and leverages its knowledge of the anonymity protocol, and the protocol specific information contained in the

messages. We quantify the relationship anonymity $P_{rel}(s,r)$ by the average probability that the attacker assigns to the messages sent by node s to node r.

Finally, we define the overhead of the anonymity network as the average path length (number of relay hops) $E[K]$ of the messages.

A. MCrowds system description

MCrowds is an anonymity network similar to Crowds [5]. The difference is that a message does not specify one node as the receiver, but it specifies a set D of nodes as receivers. The number of nodes $D = ||D||$ in the receiver set is a system parameter. For a message to reach its receiver the receiver r must be in the set D.

As in Crowds, nodes act as relays for each other. The sender initializes the receiver list D with the receiver node r and with other $D-1$ nodes chosen uniform at random. Receiver nodes D are not used as relays. The message is relayed with probability p_f, and with probability $1-p_f$ the message is sent as a multicast message to all receiver nodes D. Node r recognizes that it is the receiver while the other D-1 nodes discard the message. Note that for $D = 1$ MCrowds is equivalent to Crowds.

The mean number of hops for MCrowds is the expected value of a geometric distribution with success probability $1-p_f$ plus the multicast messages, i.e.,

$$E[K] = \frac{p_f}{1-p_f} + 1 + D \tag{1}$$

where p_f is the probability that a node will relay a message.

B. Minstrels system description

Minstrels, described in [6], uses nodes as message relays in the same way as Crowds with the difference that the number of nodes visited by a message is bounded.

When a node s wants to send a message to a node r it picks one node uniform at random among the other $N-1$ nodes and forwards the message. When a node receives a message, it checks if it is the receiver by trying to decrypt the message, or a part of it. Then the node forwards the message to another node chosen uniform at random. Note that a node does not know who is the receiver, it can only check if it is the receiver itself. The message path ends when all N nodes are visited.

To bound the path length, the messages record a list of the nodes already visited. When a relaying node receives a message, it can relay the message only to non-visited nodes. To control the maximum path length (i.e., delay) the sender can initialize the list of visited nodes with a number $M \in \{0,...,N-2\}$ of the nodes in the system. These initialized nodes are considered as visited so that the message can not be relayed to them. The sender picks the number of initialized nodes at random: it initializes the list with M nodes with probability $P(M)$, where $\sum_{M=0}^{N-2} P(M) = 1$. For $M = 0$ the list is empty, for $M = 1$ the list is initialized only with the sender and for $M > 1$ the list is initialized with the sender and $M-1$ other nodes. The list must not be initialized with the receiver, because the message would then never reach it. The distribution of $P(M)$ is a system parameter, and we use it to explore the anonymity-overhead trade-off.

The mean number of hops depends on the distribution of $P(M)$, and it can be expressed as

$$E[K] = \sum_{M=0}^{N-2} P(M)(N-M). \tag{2}$$

III. Traffic Analysis Methods and Anonymity

In the following we describe the traffic analysis methods used by the attacker to determine the traffic matrix, and outline the calculation of the relationship anonymity for MCrowds and for Minstrels for the two considered traffic analysis methods.

A. Bayesian inference method

Using this method, when the attacker intercepts a message, it considers every pair of nodes (a,b) as a possible sender-receiver pair of the message. In this case the relationship anonymity depends on two factors. First, on the probability of having an attacker node on the path, and second, on the probability that the attacker assigns to the sender (that it sent the message) and to the receiver (that it is the destination) when it gets the message. These probabilities are a function of the anonymity protocol, the number of nodes N and the number of inside attacker nodes C,

$$P_{rel}(s,r) = \sum_{i=1}^{\infty} P(\hat{S}(s),\hat{R}(r)|H_i,S(s),R(r)) \cdot P(H_i|S(s),R(r)), \tag{3}$$

where $P(H_i|S(s),R(r))$ is the probability that the position of the first attacker on the path is i given that (s,r) is the sender-receiver pair, and $P(\hat{S}(s),\hat{R}(r)|H_i,S(s),R(r))$ is the probability that the attacker identifies (s,r) as the sender-receiver pair given its position i on the path. A detailed description of calculating $P_{rel}(s,r)$ can be found in [6].

B. Maximum posteriori method

Using this method, when the attacker intercepts a message, it identifies the set Q of most likely sender-receiver pairs. The size $||Q||$ of set Q can vary from 1 (the worst case, very low anonymity) to $(N-C)\cdot(N-C-1)$ (perfect anonymity). The actual sender-receiver pair (s,r) can be either in the set $(s,r) \in Q$ or outside of it $(s,r) \notin Q$. Intuitively, we can say that $(s,r) \in Q$ is more likely than $(s,r) \notin Q$. The expression for relationship anonymity becomes

$$P_{relQ}(s,r) = \sum_{i=1}^{\infty} P(\hat{S}(s),\hat{R}(r)|(s,r) \in Q,H_i,S(s),R(r)) \cdot P((s,r) \in Q|H_i,S(s),R(r)) \cdot P(H_i|S(s),R(r)), \tag{4}$$

where $P((s,r) \in Q|H_i,S(s),R(r))$ is the probability that the sender-receiver pair is one of the most likely sender-receiver pairs, i.e., it is in the set Q, given that the first appearance of an attacker node on the path is on position i. $P(\hat{S}(s),\hat{R}(r)|(s,r) \in Q,H_i,S(s),R(r))$ is the probability that the attacker identifies (s,r) as the sender-receiver pair given the attacker node's position i on the path and $(s,r) \in Q$.

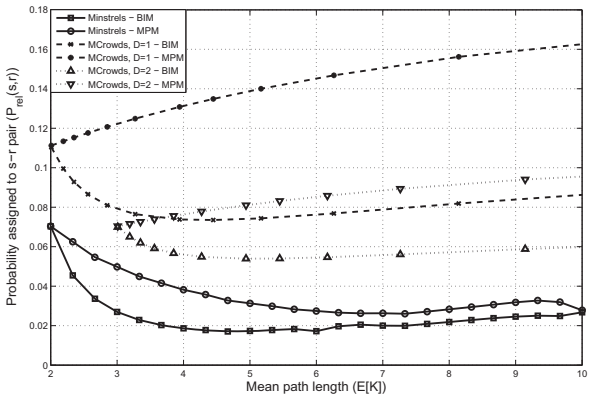

Fig. 1. Relationship anonymity vs. overhead for $N = 10$, $C = 1$.

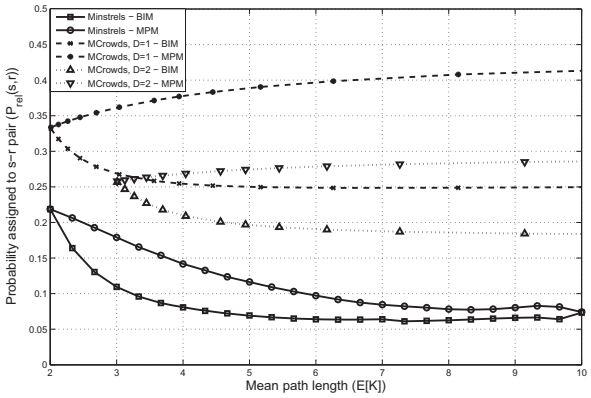

Fig. 2. Relationship anonymity vs. overhead for $N = 10$, $C = 3$.

IV. Numerical Results

To get insight into the relationship anonymity-overhead trade-off we use the two described anonymity networks, MCrowds and Minstrels, and the two attack methods, Bayesian inference method (BIM) and Maximum posteriori method (MPM). For MCrowds we use the parameters $p_f \in (0,1)$ and D of the analytical model to explore the trade-off. For Minstrels, we use various uniform, binomial, and triangular distributions for $P(M)$.

Fig. 1 shows the probability $P_{rel}(s,r)$ assigned to a sender-receiver pair as a function of the overhead (i.e., the mean path length) for $C = 1$ and $N = 10$. A higher value of $P_{rel}(s,r)$ means that the sender-receiver pair is more exposed, i.e., has less relationship anonymity. One would expect that high overhead provides good relationship anonymity (i.e., low assigned probability), but surprisingly this is not the case.

For the Bayesian inference method, above a certain level of overhead a further increase of the overhead (more relaying) has a negative effect on the anonymity for both anonymity networks. The reason is that as the number of relays increases the probability $P(H_i|S(s),R(r))$ of having an attacker node on the path increases faster than the certainty of the attacker about the identity of the sender-receiver pair decreases.

For the Maximum posteriori method, increased overhead increases the anonymity for Minstrels up to a certain level, but for MCrowds increased overhead always results in worse anonymity. We also observe that both Minstrels and MCrowds provide worse relationship anonymity for the Maximum posteriori attack method than for the Bayesian inference attack method. The reason is that the actual sender-receiver pair tends to be among the most likely sender-receiver pairs. Hence, the attacker benefits by ignoring the pairs with low probability of being the sender-receiver pair, and redistributing the assigned probability only among the most likely pairs.

Fig. 2 shows results obtained with $N = 10$ nodes and $C = 3$ attackers. Interestingly, for the Bayesian inference method relationship anonymity for Minstrels decreases above a certain level of overhead, while for Crowds the relationship anonymity improves monotonically. Hence, for $C = 3$ the probability that the attacker can assign to the sender decreases faster than the probability of having an attacker $P(H_i|S(s),R(r))$ increases.

These results lead us to two interesting conclusions. First, for an attacker it is always better to use the Maximum posteriori method than the Bayesian inference method for traffic analysis. Second, the two traffic analysis methods show similar characteristic for the relationship anonymity-overhead trade-off: best relationship anonymity might not be achieved at the highest possible overhead.

V. Conclusion and Future Work

This work is a first attempt to analyze traffic analysis methods in terms of the trade-off between relationship anonymity and communication overhead in anonymity networks. For the evaluation we considered two anonymity networks, MCrowds and Minstrels, and two attack methods. While the Maximum posteriori method always leads to lower relationship anonymity, for both traffic analysis methods the relationship anonymity is often easiest to provide at medium levels of overhead, when attackers are still unlikely to be on the path, but the sender-receiver identity is already reasonably well protected. It is subject of our future work to provide a more complete characterization of the overhead-anonymity trade-off for anonymity networks, including networks that provide probabilistic message delivery.

References

[1] D. Dzung, M. Naedele, T. V. Hoff, and M. Crevatin "Security for Industrial Communication Systems." *Proc. IEEE* vol. 82, pp. 6 1152-1177, 2005.

[2] C. W. Ten, C. C. Liu and M. Govindarasu "Vulnerability Assessment of Cybersecurity for SCADA Systems." *IEEE Trans. Power Syst.*, vol. 23, no. 4, 2008.

[3] D. Chaum "Untraceable electronic mail, return addresses and digital pseudonyms" *Commun. of the ACM 24(2)*, pp. 84-88, 1981

[4] P. Syverson, D. Goldschlag, and M. Reed "Anonymous connections and onion routing." *in Proc. IEEE Symp. on Security and Privacy*, pp. 44-54, Oakland, California, May 1997.

[5] M. Reiter and A. Rubin "Crowds: Anonymity for Web Transactions." *ACM Trans. Inform. Syst. Security*, pp. 66-92, 1998.

[6] O. Vuković, G. Dán and G. Karlsson. "On the Trade-off Between Relationship Anonymity and Communication Overhead in Anonymity Networks" *in Proc. IEEE ICC*, Jun 2011.