

Revealing Business Relationships – Eavesdropping Cross-organizational Collaboration in the Internet of Services

André Miede
Multimedia Communications
Lab (KOM) – TU Darmstadt
Rundeturmstraße 10, 64283
Darmstadt, Germany
miede@KOM.tu-
darmstadt.de

Daniel F. Abawi
Hochschule für Technik und
Wirtschaft des Saarlandes
Waldhausweg 14, 66123
Saarbrücken, Germany
abawi@htw-saarland.de

Gökhan Şimşek
Multimedia Communications
Lab (KOM) – TU Darmstadt
Rundeturmstraße 10, 64283
Darmstadt, Germany
simsek@KOM.tu-
darmstadt.de

Julian Eckert
Multimedia Communications
Lab (KOM) – TU Darmstadt
Rundeturmstraße 10, 64283
Darmstadt, Germany
eckert@KOM.tu-
darmstadt.de

Stefan Schulte
Multimedia Communications
Lab (KOM) – TU Darmstadt
Rundeturmstraße 10, 64283
Darmstadt, Germany
schulte@KOM.tu-
darmstadt.de

Ralf Steinmetz
Multimedia Communications
Lab (KOM) – TU Darmstadt
Rundeturmstraße 10, 64283
Darmstadt, Germany
steinmetz@KOM.tu-
darmstadt.de

ABSTRACT

The Internet of Services is envisioned as a global Service-oriented Architecture enabling collaboration across organizational boundaries. However, by monitoring communication endpoints, attackers can create detailed profiles of service consumers and providers even if typical security mechanisms such as message encryption are used. In a business context, this traffic analysis threatens the relationship anonymity of the participants and can reveal sensitive information about an organization's underlying business processes or a service provider's client base. In this paper, we discuss the simulation-based evaluation of different attack scenarios regarding the identification of the service compositions an organization uses. Thus, we offer insights regarding the limits of anonymity for cross-organizational collaboration in the Internet of Services.

Keywords

Security, Anonymity, Internet of Services, Service-oriented Architectures, Cross-organizational Collaboration

1. INTRODUCTION

Modern global economies have become fast-paced and highly competitive, thus, requiring organizations to adapt both quickly and continuously to changing circumstances and requirements. An important factor to achieve this goal is the

underlying enterprise Information Technology (IT), which has to integrate both internal and external systems.

The paradigm of *Service-oriented Architectures* (SOAs) [22] offers technological and organizational means in order to improve the alignment between the functional and the IT side, i. e., by enabling *service-based, cross-organizational workflows*. In the last years, *Web services* have become both a mature and successful technology for implementing the SOA paradigm.

For the near future, the *Internet of Services* is envisioned as a global SOA further facilitating cross-organizational collaboration [4,26]. The Internet of Services provides the foundation for complex business value networks by supporting the composition and aggregation of existing services to value-added services, i. e., using market places as intermediaries between service consumers and providers. Furthermore, it is a business model using the Internet as a medium for the retrieval, combination, and utilization of interoperable services. For example, market places could build compositions using services from different providers and offer these compositions as best practices for recurring process needs to service consumers.

In order to enable such service-based, cross-organizational collaboration, the security of the communication channels used, exchanged messages, and participating systems is a necessity. Regarding the security of Web service technology, substantial advancements have been achieved in the last years as discussed in the standard literature on Web service security [3,12,25]. However, several technology-independent and service-specific attacks on SOA have been identified recently, especially in the Internet of Services context [17,18].

One of these attacks aims at identifying the existence of relationships between collaborating organizations: By ob-

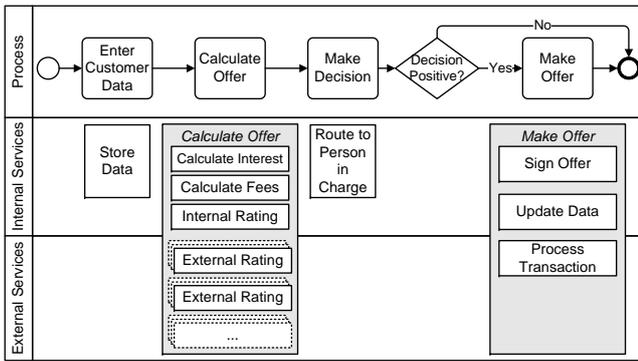


Figure 1: A fragment of a generic credit application process and possible mappings to both internal and external services.

servicing the communication between the participants in the Internet of Services, attackers can create detailed profiles of service consumers, providers, and also of market places. Because only the message exchange endpoints have to be monitored, the use of encryption or other standard mechanisms is no protection against this kind of attack, which is in general communication networks also known as “traffic analysis” [24]. In addition, due to its passive nature, it is hard to detect and depending on the monitoring means used this attack may not even be illegal. However, the obtained information reveals important organizational details, e.g., consumers exploring new business opportunities, the anticipation of mergers and acquisitions, or providers changing their business models.

The security goal that is threatened by this attack is referred to as “relationship anonymity” in the standard literature on anonymity research [23]. This means that an adversary cannot sufficiently distinguish whether the sender and recipient of a particular message are related or not. It is important to understand that this kind of anonymity does not apply to the sender and recipient of the message, i.e., they know each other. It refers only to third parties, i.e., parties that are neither sender nor recipient of the message.

A simple, but tangible example from the financial services domain is a generic credit application process, i.e., where credit ratings for customers are retrieved from an external rating agency. More about such an example is shown in Figure 1: The bank works on credit applications from its customers, e.g., first entering and storing the customer’s data in its systems (using an internal access service “Store Data”). A possible next step would be calculating the concrete credit offer, which is a service composition “Calculate Offer” consisting of both internal and external services, e.g., for external credit history ratings. After that, another service would route the resulting information to a human decider “Route to Person in Charge” for triggering the next steps. Finally, if this person’s decision is positive, the offer would be made. This results in another service composition “Make Offer”, again consisting of both internal and external services, e.g., for notifying the customer about the decision, processing the payout of the credit, etc.

However, just by monitoring the message exchange between the bank and the rating agency, an attacker can gain information about how many people apply for credit, when peak times are, when the bank works on the applications, and so on. If complete or parts of service compositions can be monitored, e.g., if information about successful credit applications is transferred to mailing and payout services via Web services, attackers could also infer a percentage of how many applications are granted or denied. This is very sensitive information rather easily available for attackers and it is not protected by the common and currently used security technology, i.e., for Web services.

A comprehensive overview of mechanisms and systems in order to achieve different types of anonymity in communication networks is given, e.g., by Edman and Yener [8]. However, even if such standard anonymity mechanisms are deployed and used correctly, attacks mounted at the edges of such networks and aiming at typical long-term business relationships are very likely to be successful. Thus, the goal of the paper at hand is the following: We investigate how an adaptation of a typical anonymity attack with respect to service compositions threatens the relationship anonymity of service consumers and providers in the Internet of Services. This is done by measuring the attacker’s success using metrics from the field of Information Retrieval while varying key system parameters such as the number of service providers, the composition complexity, or the number of observed collaborations.

The rest of the paper is structured as follows: Section 2 and 3 outline the analysis and design of our evaluation, i.e., how it was set-up and why we made certain design decisions. Subsequently, Section 4 analyzes and discusses selected results. In order to place our contributions within the body of existing research, Section 5 discusses the most relevant related work in this area. Section 6 sums up the findings and closes with a brief outlook on future work.

2. ANALYSIS

In this section, the foundation of our research is presented, i.e., the underlying assumptions, the research question, and reasons for the selected means to answer this question.

2.1 Attack Selection

First of all, how does a typical attack on anonymity and anonymity systems look like and where has it to be mounted? Assuming that basic anonymity systems are being used by the communicating participants, there are two basic choices for an attack:

1. *Attack the anonymity network itself*: The attacker tries to follow the trail of a message along the nodes of an anonymity system, e.g., as described by Guan et al. [9] (cf. Figure 2). However, this is very difficult because of the (usually) high number of participating nodes and the used security mechanisms. Thus, many nodes would have to be compromised in order to cover the whole route. In addition, a single missing node on the route makes this kind of attack even more difficult, because messages are hard to correlate between the nodes.
2. *Attack the anonymity network edges, i.e., incoming and*

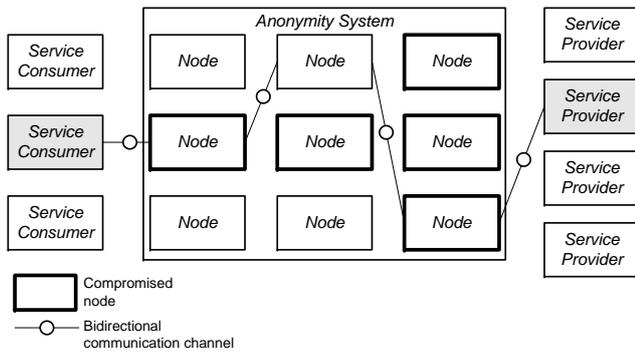


Figure 2: Attack the anonymity system itself.

outgoing messages (cf. Figure 3): This kind of attack focuses on the communication relationships of the participants and not on the anonymity network itself. These attacks are called “Intersection/Disclosure Attacks” and were introduced by Kesdogan et al. [13]. Their name is based on the intersection of the anonymity sets of senders and recipients for each communication round. However, this intersection was proven to be equal to a known NP-complete problem, thus, a statistical variant was developed by Danezis [5], reducing the required resources for the attack while still achieving good results.

Unlike the standard disclosure attacks, the so-called “Statistical Disclosure Attack” does not provide the attacker with definite information about the communication relationships but with a probability of each potential relationship. Basically, this requires the attacker to observe a large amount of interactions from which he can calculate the respective relationship probabilities. Selected details of these calculations are discussed as part of our simulation model in Section 3.2.

We chose this attack type for our evaluation, because it is a particular threat for strategic, long-term relationships, i. e., relationships that are custom in the field of service-based cross-organizational collaboration. In addition, the attack is basically independent from the used anonymity system, thus, based on certain assumptions that will be outlined below, it is a threat for most deployed anonymity systems.

For this paper, the Statistical Disclosure Attack is adapted for the Internet of Services scenario, i. e., attackers aim to identify the service compositions that organizations use for executing their processes. More details on these adaptations are given in Section 3.

2.2 Research Question

As outlined above, we assume for our research that organizations use external services (and compositions thereof) for executing their processes. Furthermore, we assume that basic countermeasures against traffic analysis are in place, thus, non-trivial attacks are needed because an attacker cannot just intercept any message in order to retrieve sender and recipient information from its header.

From this and the selected attack type follows the research question we try to answer in this paper: “How does an adap-

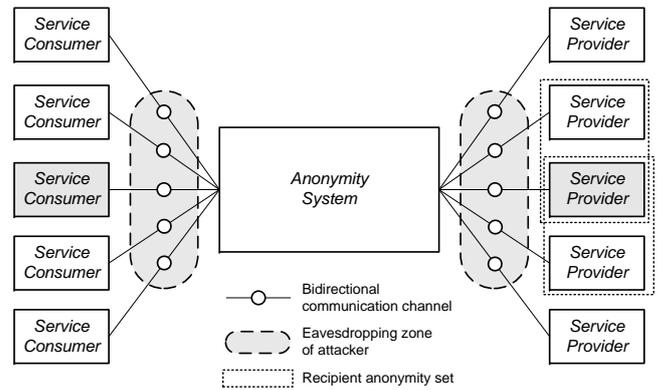


Figure 3: Attack the edges of the anonymity system.

tation of the Statistical Disclosure Attack regarding service compositions threaten the relationship anonymity of service consumers and providers in the Internet of Services?”

For answering this question, we varied key system parameters such as the number of service providers, the composition complexity, or the number of observed collaborations and measured the attacker’s success using metrics from the field of Information Retrieval. We chose a simulation-based approach because it will be an initial indicator for how well business relationships in the Internet of Services are protected against dedicated attacks. The reasons for using a simulation instead of, e. g., a testbed, are the following:

An important reason is noise reduction, i. e., dealing with specific technology (and its integration challenges) for implementing a cross-organizational testbed distracts from the attacker’s focus, which is on revealing relationships between organizations. Furthermore, the attacks need to be performed in a completely controlled environment. Our focus is not on providing a real implementation of such attacks but on developing a model that gives us information about how dangerous they are. In addition, although enabling technologies such as Web services exist, a truly global SOA such as the Internet of Services is not yet available and, thus, cannot serve as a foundation for investigating dedicated attack scenarios.

3. SIMULATION DESIGN AND SETUP

This section discusses the underlying design decisions of our simulation model, i. e., the general assumptions, an overview of the model, brief implementation information, and the different evaluated configurations are presented.

3.1 General Simulation Assumptions

For our simulation model, we assume the following regarding the different entities: The system uses end-to-end encryption that cannot be broken in time. Furthermore, it delivers messages to recipients in batches, e. g., using a so-called “Threshold Mix” [5, 13].

The attacker is passive and static, i. e., the attacker observes only and does not adapt his attack behavior. In addition, he can observe messages leaving and entering the network (not necessarily all messages) and can guess when a message

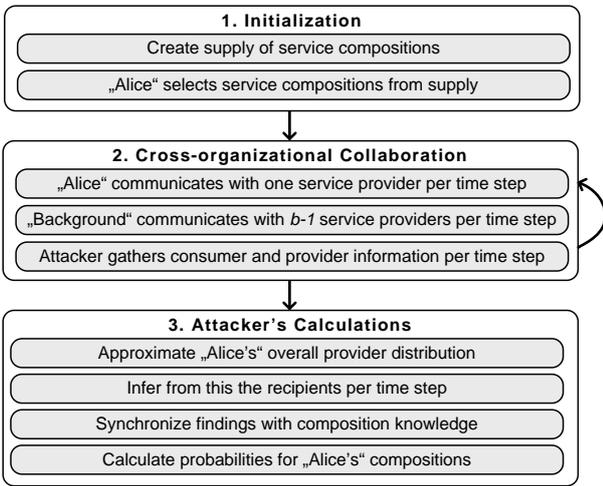


Figure 4: Schematic flow of the simulation model.

entering is likely to leave. Furthermore, the attacker knows the anonymity system’s parameters (e.g., batchsize) and the market place’s offerings, i.e., what compositions are available and what providers they consist of. Also, the participants have a consistent communication behavior, i.e., they have strategic, long-term communication relationships and do not change their service providers frequently.

These are typical and well-proven assumptions in the field of anonymity research similar to those in the related work, e.g., assuming nearly worst-case scenarios from the anonymity system’s point of view. Based on this foundation, the model is described in the next section.

3.2 Simulation Model

The attack is modeled as a stochastic model of a time-step simulation. For an overview, the basic flow of the simulation is shown in Figure 4. In addition, these steps are described in the following in more detail:

1. *Initialization*: The first step of the initialization is the generation of the overall supply of *service compositions*. Here, a service composition consists of at least one service, i.e., the ID of the respective *service provider*. The assignment of service providers to compositions is done randomly, in our model based either on a *uniform* popularity distribution of the providers or a *Zipfian* one [31]. A Zipfian distribution means that the relative probability of the i -th most popular service provider to be used is proportional to $1/i^\alpha$, leading to a more realistic selection probability of service providers. Breslau et al. showed that the requests of Web pages follow a Zipfian distribution with an exponent α of about 0.75 [2]. This finding is adopted for our simulation because it provides a realistic estimation of service offerings on the Internet. Zipfian distributions were used before in the area of anonymity research, e.g., by Shmatikov and Wang [27]. From the generated service compositions, the organization under observation, here called Alice Corp. (“Alice”) selects a certain number for executing her (business) processes.

2. *Cross-organizational Collaboration*: With the service offerings and the compositions used by Alice determined, the

cross-organizational collaboration starts. At each time-step of the simulation, Alice contacts one of the service providers that are part of her used service compositions. As in real collaboration scenarios, Alice is not the only one communicating with service providers. Thus, there is also the so-called “background”, i.e., other service consumers communicate with different service providers as well at each time-step. This background fills the remaining slots of the anonymity system’s batch of size b . The recipients of the background are denoted by the vector \vec{u} and distributed according to the general provider distribution, i.e., either uniform or Zipfian as described above. The attacker is assumed to know or approximate this distribution for his calculations. At each time-step i , the attacker intercepts the batch of messages \vec{o}_i or a fraction thereof, depending on the attacker’s spread.

3. *Attacker’s Calculations*: At regular intervals, e.g., time-step t , the attacker performs calculations for identifying Alice’s service providers in general and the corresponding service compositions in particular. The core of the calculations is based on the formal model of the classic Statistical Disclosure Attack [5]. Thus, as proven by Danezis, the attacker approximates Alice’s recipients (\vec{v}) after t time-steps based on the observed output of the anonymity system ($\sum_{i=1}^t \vec{o}_i$), the batchsize b , and the known/approximated background distribution (\vec{u}):

$$\vec{v} \approx b \frac{\sum_{i=1}^t \vec{o}_i}{t} - (b-1)\vec{u} \quad (1)$$

Using vector \vec{v} as approximated above and the stored observed vectors \vec{o}_k , the attacker calculates each vector \vec{r}_k by multiplying each element of \vec{o}_k (observed in round k) with the respective element of \vec{v} , afterwards normalizing the results using their dot product ($|\vec{v} \cdot \vec{o}_k|$):

$$\vec{r}_k = \frac{\vec{v} \cdot \vec{o}_k}{|\vec{v} \cdot \vec{o}_k|} \quad (2)$$

This then contains the probabilities about the service providers Alice communicated with in time-step k , i.e., the higher the resulting value of an element in \vec{r}_k , the more likely this service provider was used by Alice in round k .

The attacker then uses the maximum probability of each \vec{r}_k , i.e., the service provider Alice most likely communicated with in time-step k . This knowledge is then combined with the knowledge about the available service compositions, e.g., retrieved from the market places in the Internet of Services. Thus, the probability of each composition containing the most likely provider of time-step k is increased. Because the attacker does this iteratively for all observed time-steps, he builds an internal model of the service compositions Alice is using, assigning a probability to each possible composition.

The validation of the model is done as suggested by North and Macal [20]. It is based on the identified requirements, the plausibility of the assumptions, and the general development process, because the Internet of Services is not yet available for a comparison validation. These necessary aspects were discussed above and found to be valid for our model.

Using this specification as a foundation, a brief overview of

the model’s implementation is given in the next section.

3.3 Model Implementation

The simulation model is implemented using *Repast Simphony*, an agent-based modeling toolkit¹. Repast has the advantage of providing a frame for the general simulation, such as methods that are executed at each time-step of the simulation, a graphical user interface for configuring simulation parameters, and built-in functionality for tracing and logging simulation results. Furthermore, Repast models can be implemented using the Java programming language, thus, there is no need to learn yet another special modeling language.

A particular implementation aspect is the generation of random numbers based on the Zipfian (or Zeta) distribution in order to achieve more realistic results than using a basic uniform distribution for randomly selecting service providers. We calculated numbers of this distribution based on the following procedure, where $F(x)$ can be any cumulative distribution function [14]:

$$F(x) = \Pr(X \leq x), \quad y = F(x) \iff x = F^{-1}(y) \quad (3)$$

Thus, a random number X of distribution $F(x)$ can be generated by using $X = F^{-1}(U)$, where U is uniformly distributed. In our case, we used the *Apache Commons Mathematics Library* version 2.1² for calculating the inverse cumulative probability $F^{-1}(U)$ of the Zipfian distribution, extending it regarding much faster random number generation as required for our simulation runs.

For verification purposes, test cases with pre-calculated results of the attack are compared to the (non-stochastic) results of simple attack runs of our model. These test cases can be used, e. g., for verifying the results of the simulation model after changes to the underlying algorithms have been made.

The next section describes how the model and its implementation can be configured in order to reflect different attackers and attack scenarios.

3.4 Configuration

For the evaluation runs of our simulation model, it can be configured in a variety of ways, modeling different attack scenarios and attacker capabilities. The used configuration parameters are described in the following:

Service compositions are generated based on the maximum number of services per composition (mSC) and the total number of (different) service compositions (C). From these, Alice selects randomly a number of used compositions (aC).

Service providers influence the simulation by their overall number (N) and their popularity distribution, which can be either uniform or Zipfian. Furthermore, the Zipfian distribution is detailed by its skewness. Based on Breslau et al.’s seminal work on Zipfian distributions in the Web as

¹<http://repast.sourceforge.net>, last access on January 3, 2011.

²<http://commons.apache.org/math/>, last access on January 3, 2011.

shown above, we chose a skewness of $\alpha = 0.75$ for our simulations [2].

The *anonymity system* is characterized by the batchsize (b), i. e., the number of messages leaving the system per time-step.

The *attacker’s capabilities* are modeled by the parameter spread (S), which denotes the percentage of how many outgoing messages the attacker can intercept.

For the attack to have any chance of success, the following relationship must hold, as shown by Danezis [5]:

$$m < \frac{N}{b-1} \quad (4)$$

However, the parameter m , i. e., the number of Alice’s recipients, is no longer directly available in our model, because it is partially based on random variables. It can be approximated before-hand by $aC \times mSC$, which serves as an upper bound for m . At run-time, i. e., after the initialization phase, m can be determined exactly by counting the number of *distinct* service providers in all service compositions used by Alice.

As a preparation for the simulation runs, we performed a number of calibration runs in order to determine the most important parameters to be observed. The main distinction is the provider popularity, modeled by a uniform or Zipfian distribution. These are then evaluated regarding the impact of the overall number of service providers, the maximum number of services per composition, the number of compositions used by Alice, and the attacker’s spread, i. e., his access to outgoing messages. Based on these configuration decisions, the next section describes the performed simulations and discusses selected results.

4. OUTPUT ANALYSIS AND DISCUSSION

This section discusses the used evaluation metrics and selected results of the performed simulations. Due to space-constraints, some results are omitted here, e. g., the impact of the number of available compositions (C).

For each single configuration, e. g., each different value for N , 100 simulation runs were performed in order to achieve a suitable level of confidence for assessing the results [11].

4.1 Evaluation Metrics

In order to evaluate the attacker’s performance, i. e., his success regarding the identification of Alice’s service compositions, we use well-proven metrics from the field of Information Retrieval [15]. These metrics were chosen, because we consider the problem of retrieving a set of “relevant” documents from a larger set of documents to be very similar to the attacker’s goal of identifying certain compositions from the overall supply of service compositions. In addition, our scenario has the advantage of specifying definitely, what “relevant” documents, i. e., compositions, are: The ones used by Alice. Figure 5 shows the respective sets of our scenario in order to apply typical Information Retrieval metrics for our evaluation. Alice’s compositions (relevant) are denoted with A , the attacker’s identifications (retrieved) with B .

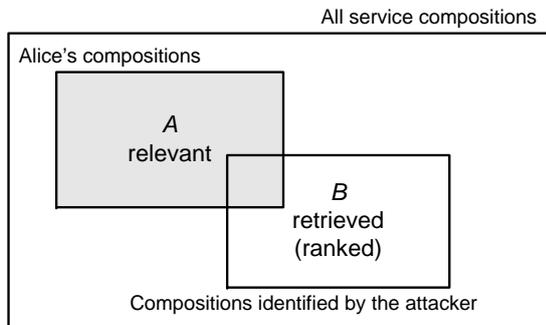


Figure 5: Retrieved/relevant sets of our attack scenario for applying Information Retrieval metrics.

Thus, the metrics mean in our scenario the following: *Precision* denotes the fraction of the identified compositions that are actually used by Alice, i. e., $\frac{|A \cap B|}{|B|}$. *Recall* denotes the fraction of how many of Alice’s compositions could be identified by the attacker, i. e., $\frac{|A \cap B|}{|A|}$.³

As the attacker assembles a ranked list of Alice’s compositions, i. e., sorted by their respective probability, this can be considered for the evaluation as well: *Mean Average Precision* (MAP) considers the position of Alice’s compositions in the attacker’s list of identified compositions. The more of Alice’s compositions are at the top of the list, thus, having a high probability, the higher the MAP. However, even if more of Alice’s compositions are identified correctly, MAP will decrease if these are ranked lower. More details on these metrics can be found in the works by Manning et al. or Mofat and Zobel [15, 19].

4.2 Impact of the Number of Service Providers

The results for different numbers of service providers, i. e., precision, recall, and MAP, are shown in Figure 6. In all these figures, the three metrics on the left are based on a uniform provider distribution while the three on the right are based on a Zipfian one. Furthermore, the y-axis uses a logarithmic scale in order to facilitate the comparison between the two distributions.

The measurements were taken after a rather short amount of interactions, i. e., 1,000 collaborations between Alice and her service providers. The reason behind this is to investigate how variations of certain system parameters influence the attacker’s results.

Uniform: Recall is basically not affected by the number of service providers, as all of Alice’s compositions are identified. However, opposed to what one might expect, precision *rises* if the number of different service providers increases. The reason for this might be, that Alice’s and the background’s interactions are distributed over a larger number of possibilities, thus, they stand out more prominently. MAP differs significantly from precision and has a very high and about constant value over the observed N . This means, that Alice’s compositions are always at the top of the list of compositions assembled by the attacker. In general, the attacker

³However, a recall of 1.0 can be achieved easily by identifying *all* available service compositions as Alice’s.

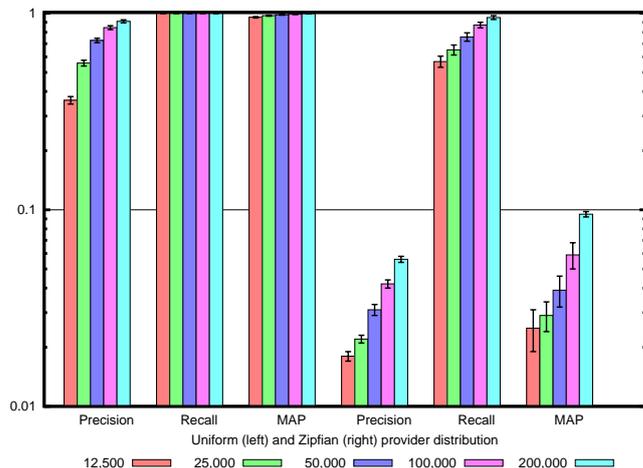


Figure 6: Varying the number of service providers N , each after 1,000 time-steps, 95% confidence intervals ($b = 125$, $C = 1,000$, $aC = 10$, $mSC = 8$, $S = 100\%$).

is very successful in this type of scenario, being mostly independent from the overall number of service providers in the system.

Zipfian: This scenario’s results are completely different for the attacker. While recall is rather good, both precision and MAP are insufficient for the attacker, achieving not even 10%. This is due to the clustering around prominent service providers, which makes it for the attacker hard to distinguish between Alice and the background. However, as seen above for the uniform distribution, the increase in providers works for the attacker, leading to more providers on the “long tail” of the Zipfian curve, thus, making profiling easier. For example, an increase by factor 4 of the service providers leads to a doubling of MAP.

4.3 Impact of the Number of Compositions Used by Alice

The results for different numbers of Alice’s service compositions are shown in Figure 7.

Uniform: The variety, i. e., number, of Alice’s compositions has a strong impact on the metrics. While the recall for 5 to 50 compositions is still very good, it declines heavily with 100 and 200 compositions used. Precision is not affected as heavily, but MAP declines even more than recall, leaving the attacker with a reasonable amount of found compositions, but which are very late in his ranked list (thus, not of much use for him). Therefore, as could be expected, Alice using more different service compositions is more difficult for the attacker.

Zipfian: As above, the attacker’s results against a Zipfian provider distribution are much worse than for the uniform one. The general trend is similar, but the degradation is more graceful than for N , e. g., recall decreases with the increase of Alice’s compositions. However, precision is rather unaffected by this increase, but for high numbers of used compositions, it even increases as well (with low and about

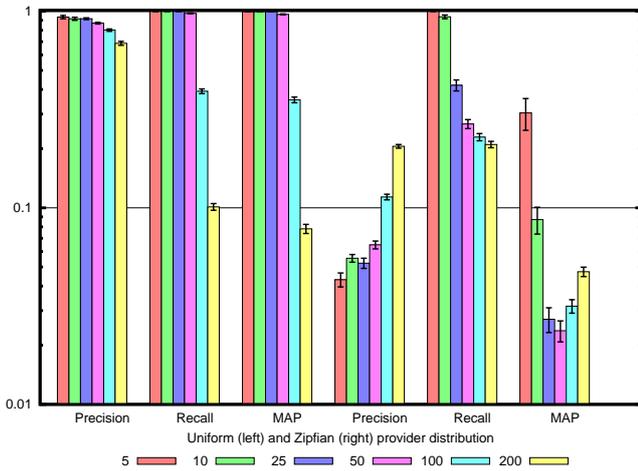


Figure 7: Varying the number of compositions used by Alice aC , each after 1,000 time-steps, 95% confidence intervals ($N = 200,000$, $b = 125$, $C = 1,000$, $mSC = 8$, $S = 100\%$).

constant values of MAP, though). This trend is due to the rising chances of the attacker identifying compositions because there are more of them.

4.4 Impact of the Maximum Number of Services per Composition

The results for different values for the maximum of services per composition are shown in Figure 8.

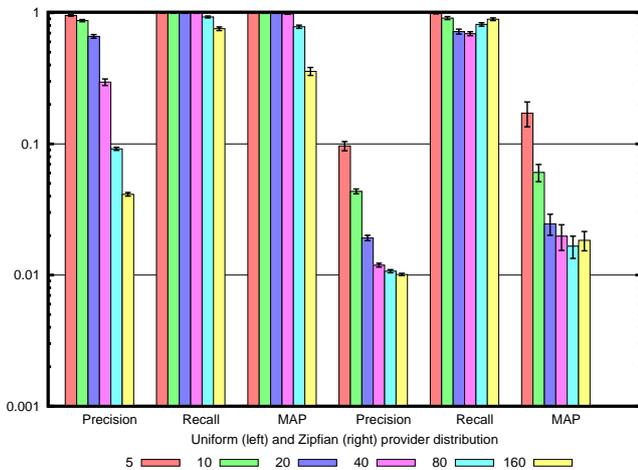


Figure 8: Varying the number of services per composition mSC , each after 1,000 time-steps, 95% confidence intervals ($N = 200,000$, $b = 125$, $C = 1,000$, $aC = 10$, $S = 100\%$).

Uniform: This parameter has a stronger impact than the number of Alice's compositions. The results are affected significantly, especially for complex compositions, i.e., high mSC between 80 and 160, which make the identification more difficult. However, recall is not affected as much, it is at most down to about 75%. The impact on MAP is not

so heavy, at least less than on precision and not as much as for high aC values. In total, this is still a success for the attacker.

Zipfian: The trend is here similar as for the uniform distribution, but the results are by far worse for the attacker. Although recall is rather high, both precision and MAP decrease significantly with a decreasing number of services per composition due to the same reasons as above. In total, this is to be regarded as insufficient for the attacker.

4.5 Impact of the Attacker's Spread

The results for different spread values of the attacker are shown in Figure 9.

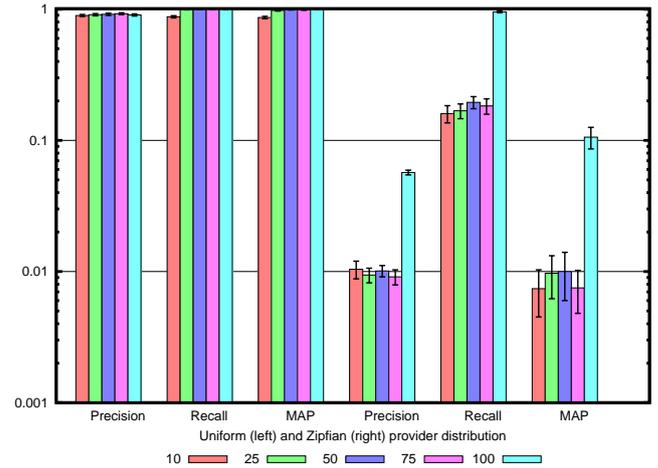


Figure 9: Varying the attacker's spread S , each after 1,000 time-steps, 95% confidence intervals ($N = 200,000$, $b = 125$, $C = 1,000$, $aC = 10$, $mSC = 8$).

Uniform: Interestingly, the spread of the attacker does not seem to influence the attacker's results very much. Except for 10% access to the anonymity system's messages, the attacker achieves very good results that are in accordance with the $N = 200,000$ runs as discussed above. A possible explanation is, that the spread only affects the time required by the attacker for achieving certain results, e.g., a 10% spread at 1,000 time-steps could lead to similar results as a 100% spread at about 100 time-steps. This aspect will be investigated further in our future work.

Zipfian: This has a devastating effect on the attacker's results. Precision and MAP fall in general below 1% (with the exception of 100% spread that is similar to the obtained results above for $N = 200,000$). Recall is better, but does not reach 20%, which is insufficient for the attacker. The presumption of the attacker needing more time if less messages can be observed should be investigated further as well in this scenario.

4.6 Impact of the Number of Time-Steps

The above evaluations considered a short, fixed amount of time in order to determine the impact of different system parameters. In addition, it is also beneficial to investigate how the metrics evolve over time, i.e., where the limits of anonymity in cross-organizational collaboration or of

the attacker could be. For this, we used the following selected scenarios from above with an underlying Zipfian service provider distribution.

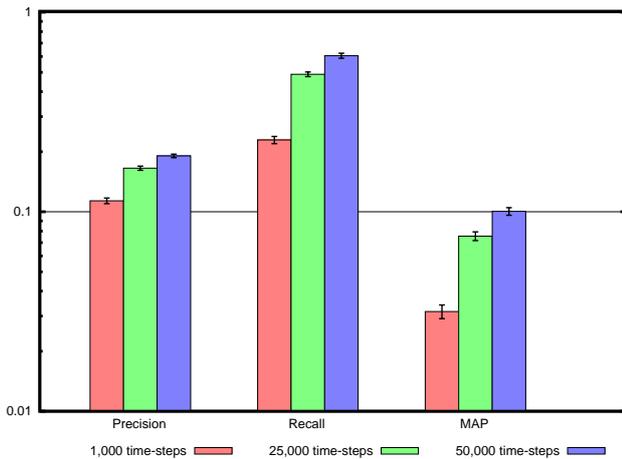


Figure 10: Evolution of attacker’s results with $aC = 100$ and Zipfian distribution, 95% confidence intervals ($N = 200,000$, $b = 125$, $C = 1,000$, $mSC = 8$, $S = 100\%$).

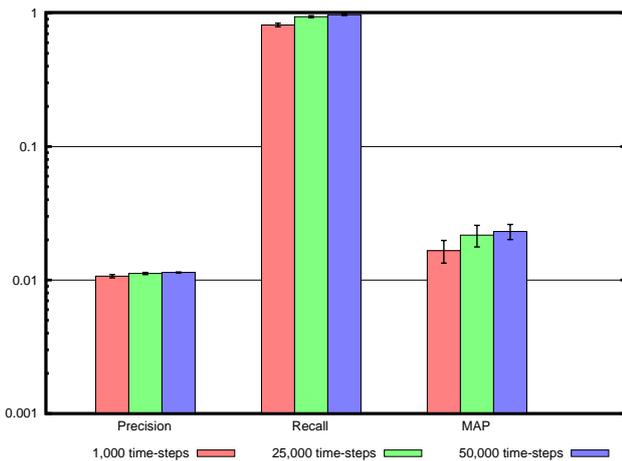


Figure 11: Evolution of attacker’s results with $mSC = 80$ and Zipfian distribution, 95% confidence intervals ($N = 200,000$, $b = 125$, $C = 1,000$, $aC = 10$, $S = 100\%$).

High number of Alice’s compositions ($ac = 100$): As shown in Figure 10, more time, i. e., more observations, gives the attacker a slight advantage. While the additional knowledge regarding MAP from 1,000 to 25,000 time-steps is high, additional 25,000 observations do not contribute much, reaching in total still only about 10%. Precision and recall do not improve much as well, so that a significant improvement after even more observations is unlikely.

High composition complexity ($mSC = 80$): Observing a scenario with a high maximum number of services per composition over a longer time does not improve the attacker’s results as depicted in Figure 11. Precision and MAP remain

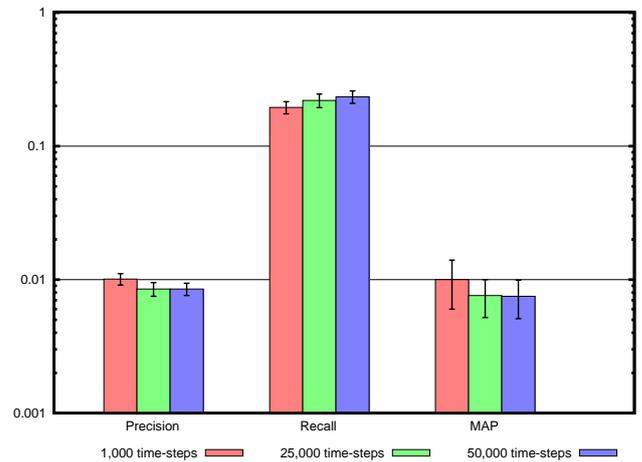


Figure 12: Evolution of attacker’s results with $S = 50\%$ and Zipfian distribution, 95% confidence intervals ($N = 200,000$, $b = 125$, $C = 1,000$, $aC = 10$, $mSC = 8$).

about constant at their very low values between 1 and 2%. However, the very high recall can be slightly improved from 1,000 to 25,000 time-steps, but not much after that.

Medium spread ($S = 50\%$): Regarding the attacker’s bad results for irregular access to the anonymity system’s messages, additional time does not help as shown in Figure 12. All observed metrics remain nearly constant at their low values, precision and MAP at most reaching 1%. As pointed out above, significant gains cannot be expected even after more time-steps.

In general, making more observations is only one approach for the attacker and not a very good one, i. e., it is most likely only used as a last resort. Other approaches, e. g., improving the internal calculations, the general model, etc. are more likely to threaten the overall anonymity in the Internet of Services. Such possibilities will be discussed below as future work.

A summary and further discussion of the overall results is given in Section 6.

5. RELATED WORK

Regarding specific attacks on anonymity, this paper focuses on attacks on the boundaries of anonymity systems, i. e., the Statistical Disclosure Attack [5] from the general class of Intersection/Disclosure Attacks [13]. Within this attack class, only simple sender-recipient-relationships were investigated so far.

However, the concept of service compositions in the Internet of Services introduces additional complexity, i. e., with respect to the relationships between service consumers, the compositions they use, and the networks of service providers that constitute these compositions. Thus, in order to gain knowledge about an organization’s processes, for example, by identifying the service compositions it uses, an attacker

has to confirm the relationships between service consumers and providers, inferring from this knowledge the service compositions this organization is most likely to use.

As a starting point for our investigations of this new scenario, we adapted the basic variant of the Statistical Disclosure Attack. Therefore, other variants and extensions were not considered so far, e.g., the use of Mix networks or different batching algorithms as described by Mathewson and Dingleline [16] or utilizing graph theory in order to relax specific user behavior assumptions of the attack model as introduced by Troncoso et al. [28].

Furthermore, the attack on relationship anonymity investigated in this paper must not be mixed up with the extensive research on Web service privacy, e.g., [10, 29, 30]. Web service privacy deals with the content of the exchanged messages, e.g., users' personal data, and how this information is further processed and possibly shared. It is an important aspect of the overall security goal "confidentiality", not of anonymity [1, 7].

On the other hand, the important aspect of anonymous communication between the different organizational participants of an SOA, i.e., with respect to third parties in order to conceal important business relationships has not been addressed so far. Further aspects of anonymity, i.e., the issue of anonymous Web service provision as well as consumption is addressed, e.g., by Papastergiou et al. [21]. However, it is questionable whether this is a desirable functionality for cross-organizational collaboration where it is important that both service consumer and provider know and trust each other, i.e., for legal reasons such as compliance or audit.

6. CONCLUSIONS AND FUTURE WORK

In this paper, we investigated the impact of attacks that aim at revealing business relationships of collaborating organizations. These attacks are of a particular danger in the field of cross-organizational service-based collaboration, because attackers can create detailed profiles of service consumers, providers, and also of market places by monitoring communication endpoints. Thus, sensitive information about the underlying business processes of the communicating organizations can be inferred easily.

Sophisticated countermeasures exist for achieving the required type of anonymity, so that an attacker cannot sufficiently distinguish whether the sender and recipient of a particular message are related or not. However, even if such standard anonymity mechanisms are deployed and used correctly, attacks mounted at the edges of such networks and aiming at typical long-term business relationships are very likely to be successful.

Therefore, this paper investigated the following research question: "How does an adaptation of the Statistical Disclosure Attack regarding service compositions threaten the relationship anonymity of service consumers and providers in the Internet of Services?"

In order to answer this question, the well-known "Statistical Disclosure Attack" was extended regarding service compositions in the Internet of Services scenario. This extension

was then evaluated using a simulation model of different attacker models and attack scenarios, which was implemented with the Repast Symphony toolkit.

While the results based on a uniform provider distribution look promising for the attacker's success, such a distribution cannot necessarily be expected in the real world, i.e., the future Internet of Services. The used Zipfian distribution, whose skew parameter is inspired by the access distribution of Web pages, has the strongest impact on the attacker's results. This leads to a clear defeat of the attacker for the observed parameters, even if more observations are made.

In addition, if the observed organization ("Alice") uses many different compositions or mainly ones with a high composition complexity, this makes the attacker's defeat even clearer. As an organization cannot just increase its process complexity for improved security, this could be achieved by extending the concept of "dummy traffic" with respect to using "dummy compositions" or "dummy services" therein, e.g., obfuscating real compositions with additional (irrelevant) services.

Further impact can be achieved by increasing the number of offered compositions or the batchsize of the anonymity system. However, increasing the batchsize is likely to have serious side-effects, e.g., regarding important Quality of Service parameters such as the response time of service requests.

These findings might suggest anonymity is not that much in danger in an Internet of Services with a suitable provider distribution. However, the attacker can also improve his chances of success by including further information such as the providers' replies into his internal calculations. In addition, outside knowledge can be used as well, e.g., a bank is more likely to collaborate with other financial service providers than with providers from the logistics, pharmaceutical, or automotive sector. These aspects will be addressed in our future work, because the threat of revealing sensitive process information remains. Thus, our next steps will be to evaluate the existing model with other configuration parameters, e.g., even more observations, and to extend the attacker model regarding industry sector information and replies from service providers, e.g., continuing the work by Danezis et al. [6].

Acknowledgments

This work is partially supported by E-Finance Lab e.V., Frankfurt am Main, Germany (www.efinancelab.de) and BearingPoint Management and Technology Consultants. Furthermore, the authors would like to thank Mr. Enkh Amgalan Ganbaatar for his support regarding the simulation runs.

7. REFERENCES

- [1] R. J. Anderson. *Security Engineering: A Guide to Building Dependable Distributed Systems*. Wiley, 2nd edition, 2008.
- [2] L. Breslau, P. Cao, L. Fan, G. Phillips, and S. Shenker. Web Caching and Zipf-like Distributions: Evidence and Implications. In *Proceedings of the Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM 1999)*, pages 126–134. IEEE, 1999.

- [3] Bundesamt für Sicherheit in der Informationstechnik. SOA-Security-Kompendium: Sicherheit in Service-orientierten Architekturen, 2009. Version 2.0. <https://www.bsi.bund.de/SOA>. Last access on January 3, 2011.
- [4] J. Cardoso, K. Voigt, and M. Winkler. Service Engineering for The Internet of Services. In *Enterprise Information Systems*, pages 15–27. Springer, 2008.
- [5] G. Danezis. Statistical Disclosure Attacks: Traffic Confirmation in Open Environments. In *Proceedings of Security and Privacy in the Age of Uncertainty, (SEC 2003)*, pages 421–426. Kluwer, 2003.
- [6] G. Danezis, C. Díaz, and C. Troncoso. Two-sided Statistical Disclosure Attack. In *Proceedings of the 7th International Conference on Privacy Enhancing Technologies (PET 2007)*, pages 30–44. Springer, 2007.
- [7] C. Eckert. *IT-Sicherheit: Konzepte – Verfahren – Protokolle*. Oldenbourg, 5th edition, 2007.
- [8] M. Edman and B. Yener. On Anonymity in an Electronic Society: A Survey of Anonymous Communication Systems. *ACM Computing Surveys (CSUR)*, 42(1):1–35, 2009.
- [9] Y. Guan, X. Fu, R. Bettati, and W. Zhao. A Quantitative Analysis of Anonymous Communications. *IEEE Transactions on Reliability*, 53(1):103–115, 2004.
- [10] P. C. K. Hung, E. Ferrari, and B. Carminati. Towards Standardized Web Services Privacy Technologies. In *Proceedings of the IEEE International Conference on Web Services (ICWS 2004)*, pages 174–181. IEEE Computer Society, 2004.
- [11] R. Jain. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley, 1991.
- [12] R. Kanneganti and P. Chodavarapu. *SOA Security*. Manning Publications, 2008.
- [13] D. Kesdogan, D. Agrawal, and S. Penz. Limits of Anonymity in Open Environments. In *Revised Papers from the 5th International Workshop on Information Hiding (IH 2002)*, pages 53–69. Springer, 2003.
- [14] D. E. Knuth. *The Art of Computer Programming, Volume 2: Seminumerical Algorithms*. Addison-Wesley Professional, 3rd edition, 1997.
- [15] C. D. Manning, P. Raghavan, and H. Schütze. *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [16] N. Mathewson and R. Dingledine. Practical Traffic Analysis: Extending and Resisting Statistical Disclosure. In *Proceedings of Privacy Enhancing Technologies Workshop (PET 2004)*, pages 17–34. Springer, 2004.
- [17] A. Miede, T. Ackermann, N. Repp, D. F. Abawi, R. Steinmetz, and P. Buxmann. Attacks on the Internet of Services – The Security Impact of Cross-organizational Service-based Collaboration. In *Proceedings of Multikonferenz Wirtschaftsinformatik (MKWI 2010)*, pages 2151–2162. Universitätsverlag Göttingen, 2010.
- [18] A. Miede, N. Nedyalkov, D. Schuller, N. Repp, and R. Steinmetz. Cross-organizational Security – The Service-oriented Difference. In *Service-Oriented Computing. ICSSOC/ServiceWave 2009 Workshops*, pages 72–81. Springer, 2010.
- [19] A. Moffat and J. Zobel. Rank-biased Precision for Measurement of Retrieval Effectiveness. *ACM Transactions on Information Systems (TOIS)*, 27(1):1–27, 2008.
- [20] M. J. North and C. M. Macal. *Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation*. Oxford University Press, 2007.
- [21] S. Papastergiou, G. Valvis, and D. Polemi. A Holistic Anonymity Framework for Web Services. In *Proceedings of the 1st International Conference on Pervasive Technologies Related to Assistive Environments (PETRA 2008)*, pages 1–8. ACM, 2008.
- [22] M. P. Papazoglou. Service-oriented Computing: Concepts, Characteristics and Directions. In *Proceedings of the Fourth International Conference on Web Information Systems Engineering (WISE 2003)*, pages 3–12, 2003.
- [23] A. Pfitzmann and M. Hansen. A Terminology for Talking about Privacy by Data Minimization: Anonymity, Unlinkability, Undetectability, Unobservability, Pseudonymity, and Identity Management, Aug. 2010. v0.34. http://dud.inf.tu-dresden.de/Anon_Terminology.shtml. Last access on January 3, 2011.
- [24] J.-F. Raymond. Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems. In *International Workshop on Designing Privacy-enhancing Technologies*, pages 10–29. Springer, 2001.
- [25] J. Rosenberg and D. Remy. *Securing Web Services with WS-Security: Demystifying WS-Security, WS-Policy, SAML, XML Signature, and XML Encryption*. Sams Publishing, 2004.
- [26] C. Schroth. The Internet of Services: Global Industrialization of Information Intensive Services. In *Proceedings of the Second IEEE International Conference on Digital Information Management (ICDIM 2007)*, pages 635–642, 2007.
- [27] V. Shmatikov and M.-H. Wang. Measuring Relationship Anonymity in Mix Networks. In *Proceedings of the 5th ACM Workshop on Privacy in the Electronic Society (WPES 2006)*, pages 59–62. ACM, 2006.
- [28] C. Troncoso, B. Gierlichs, B. Preneel, and I. Verbauwhede. Perfect Matching Disclosure Attacks. In *Proceedings of the 8th International Symposium on Privacy Enhancing Technologies (PETS 2008)*, pages 2–23. Springer, 2008.
- [29] W. Xu, V. N. Venkatakrisnan, R. Sekar, and I. V. Ramakrishnan. A Framework for Building Privacy-Conscious Composite Web Services. In *Proceedings of the IEEE International Conference on Web Services (ICWS 2006)*, pages 655–662. IEEE Computer Society, 2006.
- [30] G. Yee and L. Korba. Privacy Policy Compliance for Web Services. In *Proceedings of the IEEE International Conference on Web Services (ICWS 2004)*, pages 158–165. IEEE Computer Society, 2004.
- [31] G. K. Zipf. *The Psycho-Biology of Language: An Introduction to Dynamic Philology*. Routledge & Sons, 1999.