

Evaluating the QoS Impact of Web Service Anonymity

André Miede, Ulrich Lampe, Dieter Schuller, Julian Eckert, and Ralf Steinmetz
Multimedia Communications Lab (KOM) – Technische Universität Darmstadt
Department of Electrical Engineering & Information Technology
Rundeturmstraße 10, D-64283 Darmstadt, Germany
Corresponding Author: Andre.Miede@KOM.tu-darmstadt.de

Abstract—Web services enable the collaboration across organizational boundaries and, thus, are a powerful technology for implementing global Service-oriented Architectures, i.e., the Internet of Services. Despite typical security mechanisms such as message encryption, attackers can create detailed profiles of service consumers, providers, and market places by merely monitoring communication endpoints. 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 provider’s client base. In this paper, we evaluate the impact of using established standard anonymity mechanisms on selected Quality of Service (QoS) parameters for Web services in real networks. The obtained results aim at quantifying side-effects of using state-of-the-art countermeasures for service-specific attacks in cross-organizational collaboration.

Keywords-Web Services, Security, Anonymity, Quality of Service

I. INTRODUCTION AND MOTIVATION

In a global and highly competitive economy, modern organizations face many challenging and often rapidly changing requirements. In order to address these challenges, both the information technology (IT) side and the functional or business side have to cooperate seamlessly. In the last years, the integration of both company-wide and inter-company IT systems has emerged as a big challenge in this context.

Paradigms such as *Service-oriented Architectures* (SOAs) [1] offer technological and organizational support to improve the alignment between the business and the IT side, i.e., by enabling *service-based, cross-organizational workflows*. In the last years, *Web services* have become both a successful and mature technology to implement the SOA paradigm.

For the future, the *Internet of Services* is envisioned as a global SOA further facilitating cross-organizational collaboration [2], [3]. The Internet of Services provides the foundation for complex business 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.

In order to enable such service-based, cross-organizational collaboration, the security of the participating systems, ex-

changed messages, and used communication channels is a necessity. Regarding the security of Web service technology, substantial advancements have been achieved in the last years [4], [5]. However, in our past research, we have identified several technology-independent service-specific attacks on SOA, especially in the Internet of Services’ context [6], [7].

One of these attacks aims at identifying the relationships between collaborating organizations: By surveying the communication between the participants in the Internet of Services, attackers can create detailed profiles of service consumers, providers, and also of market places. As 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 also known as traffic analysis in general communication networks [8]. 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 details, e.g., service consumers exploring new business opportunities, the anticipation of mergers and acquisitions, or service 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 [9]. This means that an adversary should not be able to 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 outsiders, i.e., parties that are neither sender nor recipient of the message. In order to achieve different types of anonymity in communication networks, a comprehensive overview of mechanisms and systems is given, e.g., by Danezis et al. [10].

A concrete, but simple example from the financial services domain is a generic trading process, e.g., in investment banking, where market data such as interest rates and ratings are retrieved from external agencies for deal pricing calculations (cf. Figure 1). Just by monitoring the message exchange between the bank and the agencies, an attacker can gain information about the amount of requests for the internal




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

deal calculations, when the bank works on its deals, and so on. If more complex service compositions can be surveyed, e.g., if successful deals are processed by transaction services, attackers can also infer information about transactions closed successfully in general – among which are also successfully closed deals. This is easily available, but very sensitive information that is not protected by the currently used Web service security technology. However, standard anonymity mechanisms are available for communication systems which could be used until dedicated solutions are available, i.e., taking into account the high Quality of Service (QoS) requirements of Web service communication.

Because the development of new, secure anonymity mechanisms is extremely difficult and several real-life systems are available [10], the goal of this paper is to evaluate the impact of using established standard anonymity mechanisms on selected QoS parameters for Web services in real networks. The obtained results aim at quantifying side-effects of using state-of-the-art countermeasures for service-specific attacks in cross-organizational collaboration.

The rest of the paper is structured as follows: Section II discusses related work regarding Web services, QoS, and anonymity systems. After that, Section III and IV outline the preparation and design of our experiment, i.e., how it is set-up and why we make certain design decisions. Subsequently, Section V analyzes and discusses the obtained results. Section VI sums the findings up and closes with a brief outlook on future work.

II. RELATED WORK

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, which makes a discussion of related work difficult. Further aspects of anonymity such as the issue of anonymous Web service provision as well as consumption is addressed, e.g., by Papastergiou et al. [11]. 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 compliance reasons.

QoS topics, especially performance aspects such as network latency, have also been addressed previously both for Web services [12]–[15] and anonymity mechanisms [16], but not in combination. Thus, having discovered the need for relationship anonymity in Web service communication, it is the next logical step to measure its side-effects.

In addition, the attack on relationship anonymity outlined in the previous section must not be mixed up with the extensive research on Web service privacy, e.g., [17]–[19]. 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 anonymity [20], [21].

III. EXPERIMENT PREPARATION

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

In order to address the challenges of cross-organizational collaboration, we assume that Web services can be used for realizing the Internet of Services as described above. Furthermore, we assume that for protecting organizational relationships, well-proven anonymity mechanisms can be used, i.e., already deployed ones as described in the following or variations thereof.

Thus, the research question we try to answer in this paper is: “How do state-of-the-art anonymity mechanisms affect the Quality of Service of Web service-based communication?”

We choose a testbed-based approach because it will be a good indicator for the initial QoS behavior in the Internet of Services. The reasons for this are the following:

- Real networks are used, i.e., issues regarding the modeling of Internet traffic are avoided.
- Real Web services are used, i.e., unrealistic assumptions regarding self-made “dummy” Web services are avoided.
- Real anonymity systems with real users are used, i.e., modeling errors are avoided because complex systems (including user behavior) do not have to be rebuilt.

The next sections describe which QoS parameters, anonymity systems, and Web services we chose for our experiments and discuss these choices.

A. Quality of Service Parameters

QoS is an important topic for cross-organizational collaboration, because the quality of external services is crucial for the acceptance of the whole system and has to be guaranteed.




Figure 2. QoS model for service-based systems (based on [22]).

A dedicated model for QoS in service-based scenarios has been proposed by Repp [22] and is shown in Figure 2. It is selected here because it incorporates both qualitative and quantitative aspects and considers a technical point of view as well as an organizational one and, thus, goes beyond the traditional ones.

For now, we focus on an aspect of *efficiency*, i.e., the *performance* of the *service provision* (as highlighted in the figure). This is typically measured in terms of *response time*, i.e., the time it takes between sending out a service call and receiving the corresponding reply. This time consists of the network latency for the message transport and the service’s execution time on the provider side. The response time is chosen here due to its high importance for timely process execution. For the future, other QoS parameters such as throughput, reliability, and so on should be investigated as well because of their additional relevance for cross-organizational collaboration.

B. Anonymity Systems

Since the seminal work on anonymity by Chaum [23], a variety of both theoretical and live systems for providing anonymous communication has been developed and deployed [10].

One of the most famous anonymity systems is *Tor*¹ (“The Onion Router”) [16], which is chosen due to its low-latency characteristics, world-wide distribution, and easy deployment. In our experiments, *Tor* v0.2.1.26 is used.

For the *Tor* network to work, (anonymous) volunteers operate nodes from which a client chooses three at random as relays between him and a message’s destination, thus, creating a so-called circuit. Basically, a message is encrypted in an “onion-style” first using the public key of the third

node, then using the public key of the second node, and finally the one of the first node. Therefore, each node does only know the preceding and subsequent node of a message, but it does not know the whole circuit. In addition, circuits are changed about every ten minutes.

Another, well-known system is *JonDo*² as a follow-up to the *Java Anon Proxy (JAP)*, which was developed as part of the *AN.ON* project. We choose this system due to its strong security measures such as a mandatory certification for node operators, its easy deployment, and its high reputation based on the research background of its developers. Furthermore, *JonDo* has the option of using the system for free or paying for more secure access. For our experiments, we use both alternatives of *JonDo* with version 00.12.005 of the software.

Using *JonDo*, the client can choose between different “Mixes” or cascades of them. These are operated by different known providers who have to be certified in order to participate in the *JonDo* network. A *Mix* basically obfuscates the relationship between its input and output of messages so that an observer cannot link both message sets. The free and commercial versions of *JonDo* differ with respect to the available *Mix* cascades, i.e., free cascades usually consist only of one or two *Mixes* with restricted capabilities, while commercial cascades typically consist of three dedicated *Mixes*.

Last, but not least, we choose *I2P*³ (“Invisible Internet Project”) which is currently available as a beta version. However, choosing *I2P* for our experiments provides beneficial insights regarding its low-latency capabilities, e.g., in comparison with the systems mentioned above. In our experiments, *I2P* v0.7.14 is used.

More information on these systems is omitted here due to space limitations, but can be found on the respective project websites, e.g., including comparisons of different systems.⁴

C. Web Services

In order to make the experiment as realistic as possible, the Web services search engine *seekda*⁵ is used for finding real, publicly available Web services.

In total, we choose eleven different Web services from distinct and globally distributed providers as shown in Table I. The choice of the five countries at the top of the table is based on a study by Kim and Rosu [15]. This study includes information about the geographical distribution of Web services and lists these countries as top providers. Due to the large portion of the United States among Web service providers, this country contributes three providers used for this experiment. The remaining four countries in the table

¹<http://www.torproject.org>

²<http://anonymous-proxy-servers.net>

³<http://www.i2p2.de>

⁴<https://trac.torproject.org/projects/tor/wiki/TheOnionRouter/TorFAQ#ComparisontoRelatedProjects>

⁵<http://webservices.seekda.com>

Table I
GLOBAL DISTRIBUTION OF WEB SERVICE PROVIDERS USED FOR THE EXPERIMENTS.

Country	Web service provider
USA	www.webservicex.com
USA	ws.cdyne.com
USA	www.kbb.com
Australia	national.atdw.com.au
Great Britain	dw.sheetmusicdirect.com
The Netherlands	artselect.artikelbeheer.nl
Canada	netpub.cstudies.ubc.ca
Russia	www.cbr.ru
China	www.sircweb.cn
Brazil	ws.cronostelemetria.com.br
Germany	mathertel.de

are chosen in order to achieve a geographical distribution spanning the whole world.

This selection serves as a starting point only. For future experiments, the list should be extended regarding both the countries already present, i.e., new service providers, and additional countries, e.g., India.

In the following section, the design and setup of our experiment will be described based on the dimensions introduced here: QoS parameter, used anonymity system, and called Web service.

IV. EXPERIMENT DESIGN AND SETUP

This section describes, how the response time measurements are obtained in terms of the underlying measurement infrastructure. Furthermore, the experiment's typical course is discussed.

A. Measurement Infrastructure

Our measurement infrastructure is implemented as a Java-based prototype, operating fully automated on a specification of both the WSDLs and the anonymity system to be used. Results are logged in a special XML-format, facilitating post-processing and result evaluation.

For a Web service call, the URL of an WSDL-file has to be provided and from this, a suitable operation is selected. In order to anonymize the call, access to the different anonymity systems is integrated into our prototype. In addition, the configuration software of the respective anonymity system has to be running on the system performing the measurements. In order to verify the utilization of the anonymity systems for the Web service call, we conduct several tests and observe the network connections using the tool Wireshark.⁶

Then, during our experiments, we measure the elapsed time between a Web service call and the corresponding reply (response time). In addition and, e.g., for reference

measurements, it is also possible to perform pings (both ICMP and TCP), which is not used yet for our experiments.

For now, the measurements are mainly conducted locally, i.e., from different machines in Germany. Because the nodes of the used anonymity systems, e.g., Tor, are distributed globally, it is assumed that the measurements are not affected significantly by this setup, only the reference measurements that are performed not using any anonymity system.

However, in order to support this assumption, we also conduct globally distributed experiments using the Planet-Lab⁷ platform and the Tor anonymity system for comparison purposes of a local and global measurement infrastructure. For these comparisons, PlanetLab nodes in North America (Canada) are used, e.g., from the University of Saskatchewan running Tor v0.2.0.32 (due to an older operating system version on the respective machines).

The next section describes, how a typical course of our experiment is performed based on the infrastructure described here.

B. Experiment Course

Once the measurement infrastructure is implemented and verified, experiments following this course are conducted:

- 1) Input:
 - a) Choose the anonymity system that is to be evaluated, i.e., Tor, etc.
 - b) Choose the Web services to be used for the evaluation, cf. Section III-C.
 - c) Determine the number of runs, the number of Web service calls per run, and the amount of time between runs.
- 2) Experiment:
 - a) Make reference measurement, i.e., call the Web service without using a anonymity system.
 - b) Make measurement using the chosen anonymity system.
- 3) Output: The following information is gathered for each run of each Web service (and in accumulated form for all used Web services):
 - Date and time,
 - run and call number,
 - URL of the host and the used WSDL file,
 - flag for the used anonymity system the call is made with, i.e., Tor, I2P, JonDo, or NoAnon (for reference measurements not using an anonymity system),
 - response time of the Web service call, and
 - retrieval time of the WSDL file (for reference purposes).

After the overview of how the experiment is conducted and how the measurements are taken, the following section evaluates and discusses the obtained results.

⁶<http://www.wireshark.org>

⁷<http://www.planet-lab.org>, last access on September 20, 2010.




Figure 3. Overall number of measured response times per anonymity system.

V. OUTPUT ANALYSIS AND DISCUSSION

This section discusses the measurements and their implications. In total, more than 300,000 measurements were taken during different times of the day over a course of about four weeks.

A. Measurement Plausibility

First of all, in order to assess the plausibility of the measurements, their distributions are analyzed. Therefore, Figure 3 shows for each used anonymity system (and no anonymity used at all) the number of occurrences of certain response time intervals.

Regarding no anonymity system used (“NoAnon”), which means regular Web service calls, the measured values are accumulated on the left side, i.e., they have the lowest response times. Although there is a dent between 200 and 400ms, which is probably based on routing optimizations in the Internet, the overall shape and its position is distinctive and had to be expected in this or a similar appearance.

Tor, I2P, and JonDo (both free and commercial) are about normally distributed, featuring distinctive bells around 700 to 800ms (JonDoFree), 1,300 to 1,400ms (I2P and Tor), and 2,600 to 2,700ms (JonDoCom). These are plausible distributions, based on the global distribution of the systems’ nodes and their inherent random mechanisms.

The reason for the difference between the free and the commercial version of JonDo is their focus on the provided level of anonymity. While the free version uses only one or two anonymizing intermediaries, the commercial version

utilizes three of them. This explains the significant increase of the response time.

Interestingly, the reference measurements for Tor taken via PlanetLab are rather similar to the locally observed ones, they have their peak between 1,000 and 1,300 ms and show a distinctive shape as well. This can be seen as a supportive argument for the local measurements of this experiment, emphasizing the global distribution of the used systems’ nodes.

In general, the shapes of the measurements’ distributions are as expected, i.e., the reference measurements (NoAnon) are at the left-most position and the others roughly ranked by their provided levels of anonymity, with the global measurements similar to the locals ones. Thus, the measurements pass the plausibility test.

B. General Evaluation

In order to compare the different anonymity systems, i.e., their impact on the response times of Web service calls, mean response times are calculated for each Web service-anonymity system combination with 95% confidence intervals ($\alpha = 0.05$) [24] as depicted in Figure 4. A good first sign for the amount and quality of our measurements is the general small range of the confidence intervals, making comparisons easier. More details, i.e., concrete values, are listed in Table II on page 8.

As expected, using any kind of anonymity system for Web service calls increases the response time significantly, depending on the called service and used system between a factor of about 1.5 and 60.




Figure 4. Mean response times for each service per anonymity system (confidence intervals with $\alpha = 0.05$). More details, i.e., concrete values, are listed in Table II on page 8.

A ranking of the evaluated anonymity systems regarding the increase in response time seems to be directly related to the provided level of anonymity of the systems:

- 1) JonDo (free) has a mean response time between 900 and 2,200 ms, thus, features the least increase. A possible reason for this is the low number of intermediaries in comparison with the other systems. The free version of JonDo uses one or two Mixes between the sender and the destination, furthermore, if only one Mix is used, this is often one from Germany, e.g., Dresden, which can be another reason for the good performance of the overall system in the experiment. However, as there are only limited options regarding the used intermediaries, this setup cannot be changed much.
- 2) I2P and Tor perform rather similarly despite their different system setup. While I2P ranges between 1,500 and 3,000 ms, Tor is slightly slower ranging from 1,600 to 3,100 ms.
- 3) JonDo (commercial) has the highest response times as expected, due to its globally distributed three intermediary Mixes. However, the average response times still range only from 2,300 to 3,700 ms.
- 4) As described above, the distribution of the Tor measurements taken with an Canadian node of PlanetLab are rather similar to the local measurements. The slightly faster response times of the PlanetLab measurements can be explained, e.g., by the high number of North-American Web services and by the high

number of Tor nodes in this area.

Regarding the absolute increase of the response times for each service and anonymity system, the following can be observed: Comparing the absolute increase for “All” with the increase for each service, it can be seen that the mean absolute increases in milliseconds seem to be rather constant for each anonymity system: JonDo (free) at around 750 ms, I2P with about 1,500 ms, Tor around 1,700 ms, and JonDo (commercial) between 2,000 and 2,500 ms. This is an important finding as it is an approximation of the general response penalty the respective anonymity system imposes on anonymous Web service communication, irrelevant of the service providers’ geographic distribution.

With respect to the relative increase of the response times we find, that Web services with low response times, e.g., in the experiments from Germany, The Netherlands, Russia, or the USA, suffer the highest relative increase, as they rise by a factor between about 20 and 60. However, seen as whole, the increase amounts for 800 to 2,000 ms, which raises nearly all services to about the same level for each system as shown above.

Therefore, the general impact of anonymity systems on QoS for Web services, i.e., response time, can now be quantified. Thus, these experiments and their results can serve as a foundation for selecting between different anonymity systems if anonymous Web service-based collaboration is required. However, a specific recommendation depends on the context and the restrictions of the respective application

domain. Furthermore, the different levels of anonymity and underlying threat models the anonymity systems build on have to be considered as well.

VI. CONCLUSIONS AND FUTURE WORK

In this paper, we investigated the side-effects of countermeasures against a service-specific attack that threatens the security goal of (relationship) anonymity in a global SOA such as the Internet of Services.

This attack is of 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 surveying communication endpoints. Thus, sensitive information about the underlying business processes of the communicating organizations can be inferred easily.

However, 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. The achieved anonymity does not apply to sender and recipient who still know each other, e.g., for billing or audit purposes.

The impact of anonymity systems on QoS parameters is generally well-known and has led, e.g., to the development of low-latency anonymity systems. Therefore, the general research question of this paper is “How do state-of-the-art anonymity mechanisms affect the Quality of Service of Web service-based communication?”.

In order to answer this question, we prepared an experiment with a focus on measuring the QoS parameter “response time” and for this, used real, publicly available Web services, real underlying networks, and real anonymity systems such as Tor, I2P, and JonDo (both free and commercial).

The increase of the measured response times using anonymity systems seems to be directly related to the provided anonymity, i.e., ranging from JonDo (free), over I2P and Tor to JonDo (commercial). In addition, services with an already low response time receive additional times between 800 and 2,000 ms – the highest relative increase with a factor between about 20 and 60 times.

Our next steps will be to investigate other QoS parameters, such as reliability or availability, to extend the experiments regarding other anonymity systems, and to broaden the portfolio of used Web services, e.g., including services from additional countries.

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, we would like to thank Christian Gottron for his support regarding PlanetLab and both Nicolas Repp and Marco Tröbs for their valuable input.

Last, but not least, we thank the Web service providers for being part of our experiment.

REFERENCES

- [1] M. P. Papazoglou, “Service-oriented Computing: Concepts, Characteristics and Directions,” in *Proceedings of the Fourth International Conference on Web Information Systems Engineering (WISE 2003)*, 2003, pp. 3–12.
- [2] J. Cardoso, K. Voigt, and M. Winkler, “Service Engineering for The Internet of Services,” in *Enterprise Information Systems*. Springer, 2008, pp. 15–27.
- [3] 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)*, 2007, pp. 635–642.
- [4] R. Kanneganti and P. Chodavarapu, *SOA Security*. Manning Publications, 2008.
- [5] 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.
- [6] A. Miede, N. Nedyalkov, D. Schuller, N. Repp, and R. Steinmetz, “Cross-organizational Security – The Service-oriented Difference,” in *Service-Oriented Computing. IC-SOC/ServiceWave 2009 Workshops*. Springer, 2010, pp. 72–81.
- [7] 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)*. Universitätsverlag Göttingen, 2010, pp. 425–426 (short) and 2151–2162 (full).
- [8] J.-F. Raymond, “Traffic Analysis: Protocols, Attacks, Design Issues, and Open Problems,” in *International Workshop on Designing Privacy-enhancing Technologies*. Springer, 2001, pp. 10–29.
- [9] 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. Last access on September 20, 2010. [Online]. Available: http://dud.inf.tu-dresden.de/Anon_Terminology.shtml
- [10] G. Danezis, C. Díaz, and P. Syverson, “Systems for Anonymous Communication (Anonymity and Privacy),” in *CRC Handbook of Financial Cryptography and Security*, B. Rosenberg, Ed. Chapman & Hall, 2010.
- [11] 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)*. ACM, 2008, pp. 1–8.
- [12] D. Schuller, J. Eckert, A. Miede, S. Schulte, and R. Steinmetz, “QoS-Aware Service Composition for Complex Workflows,” in *Proceedings of the Fifth International Conference on Internet and Web Applications and Services (ICIW 2010)*, 2010, pp. 333–338.

- [13] D. A. Menascé, “QoS Issues in Web Services,” *IEEE Internet Computing*, vol. 6, no. 6, pp. 72–75, 2002.
- [14] S. Rosario, A. Benveniste, S. Haar, and C. Jard, “Probabilistic QoS and Soft Contracts for Transaction-Based Web Services Orchestrations,” *IEEE Transactions on Services Computing*, vol. 1, no. 4, pp. 187–200, 2008.
- [15] S. M. Kim and M. C. Rosu, “A Survey of Public Web Services,” in *Proceedings of the 13th International World Wide Web Conference (Alternate Track Papers & Posters)*. ACM, 2004, pp. 312–313.
- [16] R. Dingledine, N. Mathewson, and P. Syverson, “Tor: The Second-generation Onion Router,” in *Proceedings of the 13th Conference on USENIX Security Symposium (SSYM 2004)*. USENIX Association, 2004, pp. 303–320.
- [17] 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)*. IEEE Computer Society, 2006, pp. 655–662.
- [18] G. Yee and L. Korba, “Privacy Policy Compliance for Web Services,” in *Proceedings of the IEEE International Conference on Web Services (ICWS 2004)*. IEEE Computer Society, 2004, pp. 158–165.
- [19] 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)*. IEEE Computer Society, 2004, pp. 174–181.
- [20] R. J. Anderson, *Security Engineering: A Guide to Building Dependable Distributed Systems*, 2nd ed. Wiley, 2008.
- [21] C. Eckert, *IT-Sicherheit: Konzepte – Verfahren – Protokolle*, 5th ed. Oldenbourg, 2007, (in German).
- [22] N. Repp, *Überwachung und Steuerung dienstbasierter Architekturen – Verteilungsstrategien und deren Umsetzung*. Books on Demand, 2009, (in German).
- [23] D. L. Chaum, “Untraceable Electronic Mail, Return Addresses, and Digital Pseudonyms,” *Communications of the ACM*, vol. 24, no. 2, pp. 84–90, 1981.
- [24] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. Wiley, 1991.

Table II
MEAN RESPONSE TIMES FOR EACH SERVICE PER ANONYMITY SYSTEM (ALL VALUES IN MILLISECONDS, “CI 95” DENOTES THE 95% CONFIDENCE INTERVAL).

Web service provider	NoAnon		JonDoFree		I2P		Tor		JonDoCom		Tor-CA	
	Mean	CI 95	Mean	CI 95	Mean	CI 95	Mean	CI 95	Mean	CI 95	Mean	CI 95
www.webservicex.com	468.06	36.27	1,068.18	41.35	2,180.43	91.44	1,953.56	61.63	2,489.99	46.80	1,053.90	41.65
ws.cdyne.com	210.36	1.95	882.57	23.46	1,701.90	59.50	1,615.39	32.01	3,705.80	77.58	997.44	19.45
www.kbb.com	690.69	9.47	1,248.62	61.15	2,135.06	61.44	2,130.59	47.08	2,953.02	67.70	1,272.52	14.15
national.atdw.com.au	687.72	12.00	1,901.31	148.16	1,641.87	60.94	2,481.78	31.19	2,879.60	13.10	1,524.07	14.89
dw.sheetmusicdirect.com	1,681.54	42.08	2,146.16	65.64	3,103.21	77.48	3,076.37	64.69	3,552.82	60.36	2,349.93	54.48
artselect.artikelbeheer.nl	37.42	6.08	826.79	28.05	1,495.42	62.20	1,677.37	34.08	2,291.12	21.16	1,214.90	19.89
netpub.cstudies.ubc.ca	697.55	1.64	1,005.21	26.91	2,115.53	63.53	1,709.60	31.03	2,382.84	14.89	1,129.05	14.05
www.cbr.ru	180.79	2.77	1,459.54	45.25	1,956.58	80.55	1,755.16	35.87	2,333.17	21.63	1,384.63	28.89
www.sircweb.cn	972.32	23.72	1,579.53	55.11	2,332.91	82.67	2,690.21	77.19	2,888.96	35.44	1,681.92	31.53
ws.cronostelemetria.com.br	570.54	16.06	1,218.14	31.45	2,170.46	77.68	2,536.55	96.69	2,613.72	17.30	2,066.35	135.21
mathertel.de	38.44	1.51	895.84	41.02	1,545.49	59.17	1,448.52	28.95	2,377.59	37.85	1,175.97	15.43
All	1,029.44	17.99	1,876.32	31.76	2,488.88	30.91	2,800.96	30.24	3,305.33	32.84	1,691.40	20.58