

# e-Health for Rural Areas in Developing Countries: Lessons from the Sebokeng Experience

Massimiliano Masi<sup>1,2</sup>, Rosario Pugliese<sup>1</sup>, and Francesco Tiezzi<sup>3</sup>

<sup>1</sup> Università degli Studi di Firenze, Viale Morgagni 65, 50134, Firenze, Italy

<sup>2</sup> Tiani “Spirit” GmbH, Guglgasse 6, Gasometer A, 1110, Vienna, Austria

<sup>3</sup> IMT Advanced Studies Lucca, Piazza S. Ponziano 6, 55100, Lucca, Italy

massimiliano.masi@tiani-spirit.com, rosario.pugliese@unifi.it  
francesco.tiezzi@imtlucca.it

**Abstract.** We report the experience gained in an e-Health project in the Gauteng province, in South Africa. A Proof-of-Concept of the project has been already installed in 3 clinics in the Sebokeng township. The project is now going to be applied to 300 clinics in the whole province. This extension of the Proof-of-Concept can however give rise to security flaws because of the inclusion of rural areas with unreliable Internet connection. We address this problem and propose a safe solution.

**Key words:** e-Health systems in developing countries, information security, healthcare technology standards, e-Health experiences

## 1 Introduction

In recent years the importance of healthcare systems based on Electronic Health Records (EHRs) has been addressed by governments and institutions. An EHR is a set of sensitive data written in a machine readable format (e.g., the HL7's CDA [1]) containing the healthcare history of a patient, such as vital signs, prescribed medicines, billing, and the patient summary.

Many projects have been started worldwide with the aim of developing systems for electronic healthcare (e-Health) based on EHRs sharing and capable of providing optimum patient care (see, e.g., [2, 3]). To ensure fallbacks and interoperability, such systems need to be built using international, well-known, and open standards, due to the impact on financing and governance that they might have. Large investments are conducted by governments, therefore a good confidence in the project success is needed. For this reason the initiative Integrating the Healthcare Enterprises (IHE, [4]) was founded with the goal of tailoring already existing standards (such as [1, 5, 6]) in the context of Service-Oriented Architectures for setting up e-Health projects. IHE thus provides a standard methodology for building applications in which EHRs exchange plays a significant role.

By adopting standard technologies, governments and hospitals obtain significant benefits from both a financial and a practical point of view. E-HR.GP [7] is an example of an ongoing project based on IHE standard technologies aiming at managing EHRs of patients from the Gauteng province, South Africa. A Proof-of-Concept (PoC) of such project has been successfully installed in the Sebokeng township [8]. The next step of the E-HR.GP project is the extension of the PoC implementation to cover the clinics of the overall Gauteng province.

The process of creating e-Health systems, however, is not just a merely adoption of healthcare standards. Indeed, these standards are often designed with the assumption of a high-speed network infrastructure in place, which is not always the case when considering rural areas of developing countries. This can result in a set of security flaws in the authentication process of healthcare professionals due to missing requirements on technology (e.g., implicit assumptions on communication channels, such as availability and reliability). Therefore, since a considerable area of the Gauteng province suffers from the above mentioned connectivity problems, we believe that such security flaws would occur in the extension of the PoC, unless suitable measures are taken. In this paper, we tackle the problem of amending the system developed in the E-HR.GP's PoC by proposing the use of a formally-proved correct IHE-based protocol [9], so to include also clinics in rural areas with low or absent Internet connectivity and reachable only by hundred of kilometers of sand tracks.

To sum up, the main contributions of this work are: 1) the presentation of the outcome of the current status of the E-HR.GP's PoC; 2) the identification of potential authentication problems that can arise in the extension of the PoC; 3) the proposal of adopting of a formally-proved secure protocol for connecting clinics in rural environments.

The rest of the paper is structured as follows. In Section 2, we provide an overview of the E-HR.GP project. In Section 3, we discuss major problems to tackle when extending the application of the project to include more clinics and how to solve them by means of an IHE-based protocol. In Section 4, we touch upon more closely related work. Finally, in Section 5, we conclude the paper.

## 2 The E-HR.GP project

Before the starting of the E-HR.GP project, some clinics of the Gauteng region were already equipped with Clinical Information Systems [7]. These systems, however, were not interoperable and were not based on international standards, resulting in an high cost of ownership. In fact, in each clinic, long patient queues and slow response times were experienced, which often led to preventable deaths.

To provide a better patient care, in 2007 the Gauteng Department of Health (GDoH) issued a request for proposal (RFP) for enabling e-Health in the region. Fourteen companies answered to the RFP, and six of them were shortlisted. At the end, one preferred supplier and one reserve were chosen. The preferred sup-

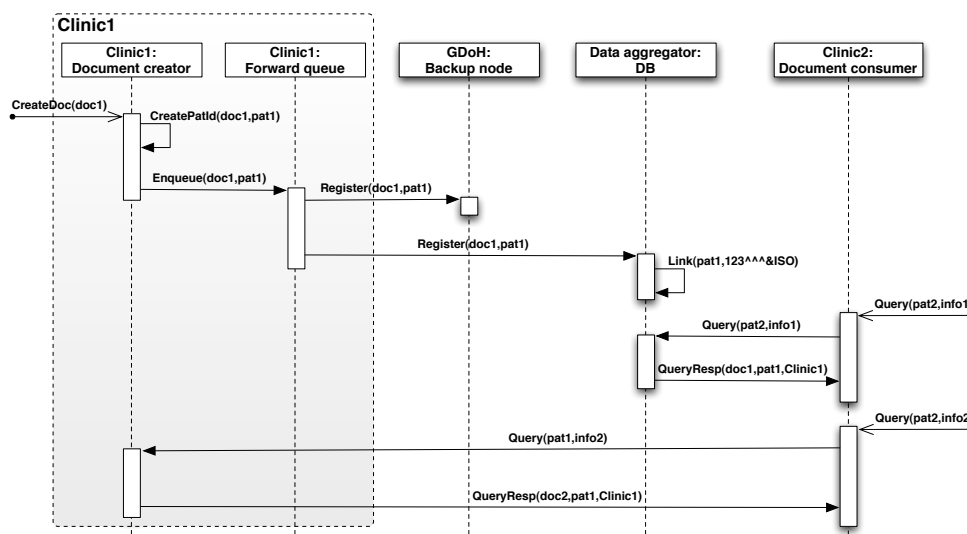


Fig. 1. Messages exchange in the E-HR.GP Proof-of-Concept

plier (the Baaki Consortium<sup>1</sup>) realized the PoC within the time frame defined (from august 2007 to august 2008). The project is divided in three phases: the first phase consists of the definition of the PoC, the second is the full implementation of the PoC, and the third the application to the whole region. In the first two phases, three clinics (Johan Deo, Levai Mbatha, and Dr. Helga Kuhn Clinic) and the 700-beds Sebokeng hospital were involved.

### 2.1 Proof-of-Concept software architecture

The three clinics and the hospital run the same software components (actually, the hospital is considered as a ‘big’ clinic). The common messages exchange in such an architecture is depicted in Figure 1, where a clinic creates a document belonging to an EHR, which is spread over the clinic’s network, and then another clinic tries to fetch it.

Specifically, a document *doc1*, e.g., a patient summary, is created (and locally stored) in a generic clinic, *Clinic1*. The document belongs to a patient, whose identifier *pat1* is chosen by *Clinic1*. The document, its metadata and the newly created patient identifier are enveloped in a message queued in a software actor, named *Forward queue*. When the *Forward queue* processes the message, this is duplicated: one copy is sent through the IHE-based transaction *Register* to a node located into the GDoH offices, while using the same transaction a second copy is sent to the *Data aggregator*. The *Backup node* installed at the GDoH offices has the task to provide a working backup image for each clinic. This reliable node should contain the data of all clinics. The *Data aggregator* is another reliable

<sup>1</sup> Baaki is a consortium created by two South African resellers, Equiton Investment (Pty) Ltd and AMETHST (Pty) Ltd.

node that instead acts as a forwarder: whenever a request for a given patient is received, it returns its copy of the document and the link to the clinic where the data were originally stored; every subsequent query about the patient will be delivered directly to the clinic, without passing through the **Data aggregator**.

Notably, the **Data aggregator** contains all the patients' local identifiers and their data as registered in the clinics. These local identifiers are linked together (through a **Link** transaction) by means of an identifier that is unique for the whole Gauteng province. In our example **pat1** is linked to 123^^^&ISO (the patient identifier is written using a notation *à la* H17 [1]).


Now, suppose that a doctor sitting in **Clinic2** wants to retrieve a document containing information **info1** for a patient identified by **pat2**. **Clinic2** does not know where information about the patient **pat2** can be retrieved, therefore the **Data aggregator** is contacted, using the IHE-based transaction named **Query**. The **Data aggregator** knows the patient identifier, which is actually linked to 123^^^&ISO too (i.e. **pat1** and **pat2** identify the same patient, resolved by means of *demographics queries* [4]). This patient has documents into **Clinic1** and, hence, the **Data aggregator** returns this information to **Clinic2**, together with the patient identifier **pat1**. For a given amount of time, this information is cached in **Clinic2** and all the subsequent queries for **pat2** will be sent to **Clinic1** directly. Indeed, in the last pair of **Query/QueryResp** transactions in Figure 1, a document consumer asks for information **info2** for patient **pat2**; **Clinic2** now is informed that the patient is known as **pat1** at **Clinic1** and, hence, directly sends the query to **Clinic1** for patient **pat1**.

## 2.2 The IHE model

Each clinic involved in the PoC uses an e-Health system based on the IHE Cross Enterprise Document Sharing (XDS) model [4], depicted in Figure 2. The XDS model uses a central document registry that acts as a catalogue for the data. The document source (e.g., a medical device) produces healthcare data for patients and store data into one or more document repositories (e.g., databases). The repositories extract metadata and update the registry. Then, possibly quite long time later, the document consumer (e.g., a doctor's workstation) queries the registry and obtains a link to the repositories from where data can be downloaded and displayed to the doctor.

In IHE standards, building blocks of e-Health systems are defined in *profiles* that expose a set of requirements. Two or more profiles are *grouped* together by merging their respective requirements. The resulting system can be grouped again with other profiles for building a complete e-Health solution.

When dealing with such sensible data as EHRs, security and privacy play a crucial role. Aspects like confidentiality, authentication, integrity and authorization are crucial for the success of any e-Health project. The IHE security model is based on the Audit Trail and Node Authentication (ATNA) profile [4]: each system is classified as either *secure application* or *secure node*. A secure application is a system that permits to establish Transport Layer Security (TLS) channels for exchanging medical records using IHE-defined transactions. A secure node



**Fig. 2.** The XDS model grouped with ATNA and XUA profiles

is a secure application where no other way to access patient's data exists (i.e. there is no physical access to the machine). With these requirements, each actor (e.g. a document consumer) is authenticated by exchanging X509 certificates and corresponding private keys. Instead, healthcare professional authentication is defined in the Cross Enterprise Document Assertion (XUA) profile [4]. If a system is grouped with this profile, each message must contain a Security Assertion Markup Language (SAML) assertion [6], i.e., a security token encoded using the XML language, providing the digital identity of the professional and issued by a trusted third party named *identity provider*.

Creating a system grouping standard IHE profiles permits to drastically reduce errors due to missing healthcare information in paper-based clinics. EHRs are always available to every patient and authorized healthcare professional across the clinics. The probability to perform, e.g., duplicate medical tests and examinations is then reduced significantly. The hardware running in each clinic is reduced in costs: with an investment of a few thousands of U.S. dollars for off-the-shelf components, a clinic can start to be operational in a few days.

### 3 Applying the protocol for disconnected clinics

The roadmap of the programme after the PoC phases is to increase the number of hospitals and clinics involved in the E-HR.GP project. Although the Gauteng

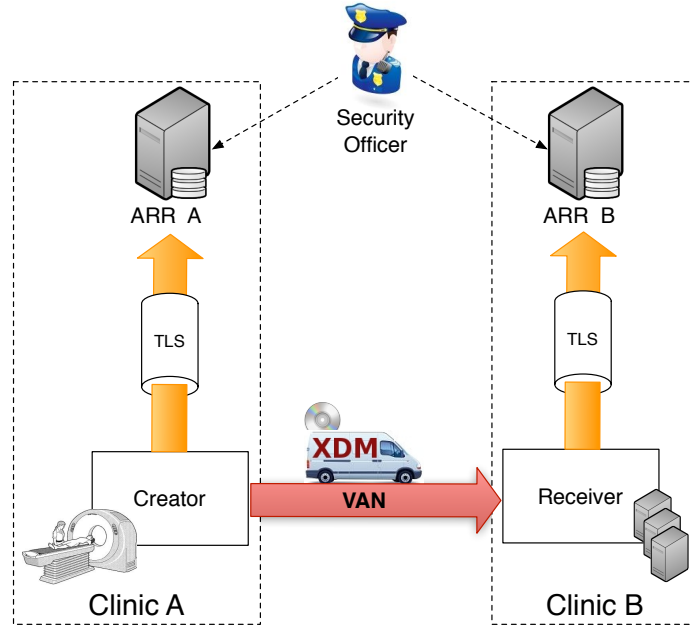
province is the geographically smallest province in South Africa, nowadays it is the most densely populated<sup>2</sup>. Thus, up to 38 hospitals, 30 community health clinics and around 300 clinics will participate in the exchange of medical records.

The PoC setting described in Section 2 fits perfectly within regions with a high-speed Internet connection. Indeed, healthcare documents usually have important dimensions (e.g., the size of DICOM [5] images can be larger than one gigabyte) and, thus, their real-time exchange is expensive. For this reason, nightly transfers to the central repository (i.e. the backup node) of the GDoH are scheduled. Each transaction passes through a forward queue (see Section 2.1) that duplicates the messages: a copy is sent to the local XDS affinity domain, while the duplicated message is queued waiting to be transferred when the connection link is operating. Unfortunately, in Gauteng, due to the limitations of the Internet connectivity in some spots of the region (where network connection is mainly based on GPRS and 56k modems), nightly transfers among clinics and the repositories may not reach the termination, due to the large size of files.

The IHE initiative solves the above problem by introducing the Cross Enterprise Document Sharing using Portable Media (XDM) profile [4]. This specification enables the usage of different types of electronic support, like CDs, DVDs, and USB drivers, for the transmission of EHRs. By using the IHE grouping mechanism mentioned in Section 2.2, XDM replaces in the XDS model the TLS channels established from the forward queue to the endpoint of Data aggregator and Backup node with out-of-band channels, based on a car transportation system. However, simply adopting the newly created system is not sufficient. In fact a new security flaw would be introduced: by dropping the TLS channel, the system would miss the authentication of the remote peer (the clinic). The receiving system would not be able to trust the message carried by the portable media because no key material is present that enforces authenticity. This flaw would introduce serious drawbacks: a malicious intermediary could change the message on the channel (resulting in a compromise of the safety of the patient) or could replay messages to get more medicines prescribed (e.g., for the illegal market). In this way, the system security would be totally demanded to the security of the out-of-band channel. Usually this channel is represented by a van travelling among clinics for collecting portable media to be stored in the central repository. The van is then a critical actor in the project security model.

A possible solution of such an authentication flaw, which does not put any requirement on the security of the van itself, is presented in [9]. This work introduces an ad-hoc protocol for exchanging medical records that exploits the XDM standard. The protocol does not affect software already in production stage and its adoption has a reduced impact on costs. Security of the protocol is proved by using the formal methods described in [10], in particular the process calculus COWS for formally modeling the protocol, the temporal logic SoCL for expressing the protocol's security properties, and the model checker CMC for verifying the properties over the model.

<sup>2</sup> Around 9 millions, of which 90% are urbanized.



**Fig. 3.** The protocol for sharing medical records in disconnected clinics

The protocol is depicted in Figure 3. Each clinic has its own information system where the communication channel between the actors **Creator** and **Receiver** and the corresponding *Audit Record Repositories (ARRs)*<sup>3</sup> is assumed to be secure. Indeed, a clinic in a rural area is usually built in an environment composed of a single machine or a few machines running in the same private local network disconnected from the Internet, where network intrusions cannot be performed. On the other hand, if a clinic is connected to the Internet (e.g., an hospital), software like Intrusion Detection Systems (IDS) [11] can assure the authenticity of the above mentioned channel. The scenario in Figure 3 is common for clinics in rural areas, in which the information systems are running on ATNA secure nodes. XDM, ATNA and XUA are grouped together, therefore the security token that authenticates each message is a SAML assertion as per the XUA profile.

The protocol is based on an end-to-end scenario, where clinics cooperate together to exchange messages. In the communication scenario, the clinic that produces a document is named **Creator** and the clinic that receives the document is named **Receiver**. Each clinic contains a security token service that is able to authenticate the healthcare professional, and an ARR, as part of ATNA. The ARR is a tamperproof hardware that stores the logging entries (i.e., the audit trails) for all the transactions performed. The software components establish a TLS channel with the ARR, as per ATNA specifications [4]. Whenever a document is created, or a document needs to be retrieved, the corresponding XDS

<sup>3</sup> The role of ARR A and ARR B, as well as the role of the Security Officer, will be explained later on.

message is enqueued into the forward queue. When the forward queue needs to be transferred, the XDM component writes the messages to the portable media. This media is transferred to the receiver along the ‘communication channel’ represented by the van. The protocol, indeed, abstracts this car-transportation system as a communication channel that permits to send a single chunk of data. In fact, this channel does not satisfy any mutual authentication properties [12]. Each message contains a cryptographic secret created by the identity provider of the creator clinic and encrypted for the clinic acting as a receiver. The secret is also stored in the ARR of both clinics and it is the key used for tracing the transaction and the message exchange in the channel represented by the van. No information about the van or its driver needs to be stored in the messages or in the associated audit trails. The message structure is maintained as per the XDS/XDM specifications.

In order to perform analysis on the security of the protocol, a threat model is proposed in [9]. In such a model, an attacker *à la* Dolev-Yao [13] is used. Based on the past experience of other e-Health projects (see, e.g., the NHIN threat model [14]), our threat model identifies four different attacks with a particular attention to the scenario in rural areas. Most of the attacks are performed (sometimes physically) to the car-transportation system. In the first attack, the intruder suppresses a message (e.g., a submission for a new rare form of allergy). In the second attack, the intruder is trying to impersonate an healthcare professional by reusing a token found in a stolen portable media. In the third attack, the payload of the message is manipulated so to reuse the authentication token for obtaining different resources (e.g., different prescribed medicines). In the fourth attack the intruder tries to replicate the message for obtaining the same resource multiple times.

All the attacks are discovered by the design of the protocol, except the first. In fact, in the scenario under study, where assuming clinics’ clocks be synchronized is a too strong requirement, a new actor, called *security officer*, is introduced to discover the threat. This officer, that can be enacted by a human being polling the clinics in a round-robin fashion, reads the audit record repository to keep track of the logging entries. By performing this action, he/she recognizes time differences between the message sent and the message received by the clinic nodes and is thus able to discover attacks based on message suppression. Notably, the security officer is assumed to be a trustworthy actor. This assumption is reasonable in the considered scenario. In fact, the security officer can only manage signed and encrypted data and, moreover, a suppression of logging entries is detected by the protocol.

It is worth noticing that the protocol can scale with a large number of clinics. Indeed, we can apply the XDM-based protocol between disconnected clinics and the classic XDS model in the other cases, and then exploit the hierarchical organization of clinics enabled by XDS. The protocol is also designed to scale with a great number of patients, by relying on the fact that a van can carry many portable media and that, nowadays, electronic supports like DVDs and USB drivers can store significant amounts of data.



## 4 Related work

Providing EHRs for rural areas and disconnected clinics is a challenging research field. An approach using Mesh networks is proposed in [15]. This approach has a great benefit: a Mesh network infrastructure, which can be exploited for e-Health purposes, is implementable with low-cost hardware. Our approach instead acts at the software infrastructure level rather than at the hardware one, and has basically no additional costs, except those for adapting the software infrastructure to be able to fulfill the requirements of [9] (to register secrets into ARRs and to copy authentication tokens into portable media). If detection of lost messages is needed, a security officer must be employed and trained.

Our protocol, firstly proposed in [9], has the advantage to be formally proved against errors in a specific threat model. Formal analysis on security protocols has given very good results in the last years (see, e.g., [16, 17, 18, 19]). A more complete survey can be found in [20, 21]. The work closest to ours is [22] where a generic communication infrastructure is proposed within a governmental project in Mozambique. In that work, a “fat” client approach is taken, in which the application and persistency layer reside on the local machine. Differently from our approach, data exchanged are partially anonymized, since they are used for statistical purposes, while we adopt a protocol that ensures authentication and confidentiality of sensitive data.

## 5 Conclusions

In this paper we reported the experience on the first phases of the project EHR.GP running in Sebokeng, Gauteng, a province of South Africa. This project aims at giving optimum patient care in the province by enabling an e-Health solution based on international standards. The major drawback of this and other similar projects in developing countries with rural areas is the low or absent network connectivity. To cope with problems due to the lack of connectivity, we propose to exchange messages by exploiting the international IHE standard XDM, which permits the usage of portable media to share EHRs. However, adopting this standard drops the requirement of having authenticated channels. To recover authentication, we propose the use of a formally-proved secure protocol that amends the standard in order to provide a secure EHRs exchange with disconnected clinics using a car-transportation system to bear portable media.

As we show in this work, just adopting international standards for the establishment of e-Health projects is not sufficient, because new security flaws can be introduced. Our aim is to continue to use formal methods to analyse e-Health projects for developing countries in order to tailor secure solutions with an important level of technology and reduced costs.

## References

1. Health Level Seven organization: HL7 standards (2009) <http://www.hl7.org>.

2. The epSOS project: An European eHealth Project (2010) <http://www.epsos.eu>.
3. The Nationwide Health Information Network (NHIN): An American eHealth Project (2009) <http://healthit.hhs.gov/portal/server.pt>, last visited on 4 July 2011.
4. The IHE Initiative: IT Infrastructure Technical Framework (2009) <http://www.ihe.net>.
5. ACR-NEMA: Digital Imaging and Communications in Medicine (DICOM) (1995)
6. OASIS Security Services TC: Assertions and protocols for the OASIS security assertion markup language (SAML) v2.02 (2005) <http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf>.
7. M. Mosupi, Gauteng Department of Health (GDoH): E-HR.GP, System Implementation (2008) [http://www.ihe-austria.at/fileadmin/user\\_upload/CAT2009/documents/MmakgosiMosupiIHE\\_Conference\\_Presentation\\_220409.pdf](http://www.ihe-austria.at/fileadmin/user_upload/CAT2009/documents/MmakgosiMosupiIHE_Conference_Presentation_220409.pdf) (last visited on 4 July 2011).
8. Wikipedia, the free encyclopedia: Sebokeng, Gauteng. [http://en.wikipedia.org/wiki/Sebokeng,\\_Gauteng](http://en.wikipedia.org/wiki/Sebokeng,_Gauteng) Wikipedia entry (last visited on 4 July 2011).
9. Masi, M., Pugliese, R., Tiezzi, F.: Security analysis of standard- driven communication protocols for healthcare scenarios. In: Healthcom, IEEE (2011) 308–315
10. Fantechi, A., Gnesi, S., Lapadula, A., Mazzanti, F., Pugliese, R., Tiezzi, F.: A Logical Verification Methodology for Service-Oriented Computing. *ACM Transactions on Software Engineering and Methodology* (2011) To appear.
11. Wikipedia, the free encyclopedia: Intrusion Detection Systems. [http://en.wikipedia.org/wiki/Intrusion\\_detection\\_system](http://en.wikipedia.org/wiki/Intrusion_detection_system) Wikipedia entry (last visited on 4 July 2011).
12. Lowe, G.: A Hierarchy of Authentication Specifications. In: CSFW, IEEE (1997) 31–44
13. Dolev, D., Yao, A.: On the security of public key protocols. *IEEE Transactions on Information Theory* **29**(2) (1983) 198–207
14. The Nationwide Health Information Network (NHIN): Threat models (2009) <http://wiki.directproject.org/Threat+Models>, (last visited on 4 July 2011).
15. Yarali, A., Ahsant, B., Rahman, S.: Wireless mesh networking: A key solution for emergency & rural applications. In: MESH, IEEE (2009) 143–149
16. Armando, A., Carbone, R., Compagna, L., Cuellar, J., Abad, L.: Formal Analysis of SAML 2.0 Web Browser Single Sign-On: Breaking the SAML-based Single Sign-On for Google Apps. In: FMSE, ACM (2008) 1–10
17. Johnson, J., Langworthy, D., Lamport, L., Vogt, F.: Formal specification of a web services protocol. In: WSFM. Volume 105 of ENTCS., Elsevier (2004) 147–158
18. Blanchet, B.: CryptoVerif: Computationally Sound Mechanized Prover for Cryptographic Protocols. Dagstuhl seminar “Formal Protocol Verification Applied” (2007)
19. Bhargavan, K., Corin, R., Fournet, C., Gordon, A.: Secure sessions for web services. In: SWS, ACM (2004) 56–66
20. Ma, L., Tsai, J.: Formal verification techniques for computer communication security protocols. *Handbook of Software Engineering and Knowledge Engineering* **1** (2001) 23–46
21. Fidge, C.: A Survey of Verification Techniques for Security Protocols. Technical Report 01-22, Software Verification Research Centre, Univ. of Queensland (2001)
22. Armellin, G., Bogoni, L., Chiasera, A., Toai, T., Zanella, G.: Enabling business intelligence functions over loosely coupled environment. In: AFRICOMM, LNICST, Springer (2010)