Collaborative Fake Media Detection in a Trust-Aware Real-Time Distribution Network

Dominik Renzel, Khaled A. N. Rashed, Ralf Klamma

Informatik 5 (Information Systems & Databases), RWTH Aachen University
Ahornstr. 55, D-52056, Aachen, Germany
{renzel,rashed,klamma}@dbis.rwth-aachen.de

Abstract. Due to the increased incorporation of external sources media agencies face the challenge of providing high-trust media to their customers. Automatic image processing approaches still do not bridge the semantic gap to identify fakes. Complementary community-based approaches lack real-time media distribution for improved awareness and base trust on subjective opinions instead of objective actions. In this paper we propose a collaborative fake media detection approach addressing these challenges in form of a federated, trust-aware media distribution network. Starting from a realistic use case scenario we elicit requirements and present an XMPP-based and Web service-enhanced multimedia distribution network as solution. Finally, we sketch a Web-based fake media detection application powered by our network and its services.

1 Introduction

Traditionally, people consider images as a means for true reproduction of real events and accepted as a proof of occurrence of such events. Recently, this consideration is not longer valid since fake images have a high occurrence especially now that images can be faked and distributed arbitrarily without much effort. Nowadays, news creation processes have taken significant distance from being conducted in isolation. Following the basic principles of the Open Innovation approach [3], in today's media distribution networks different communities are involved as both information providers and consumers. With the growing availability of low-cost high-quality multimedia processing and context sensor equipment in mobile devices, it has already become widespread practice to even have amateur reporters on site of interesting events serve as information sources. With such an inherently distributed approach, the authenticity of distributed multimedia is even more endangered than in previous more isolated approaches. Today's media thus face the challenge of deciding if media are real or faked, ideally before they are further broadcasted to their customers, who pay for high-trust media.

Consider the following infamous cases where faked media were finally published to information end-consumers. A recent example of a faked image manipulated by the newspaper Al-Ahram and published in international media is showing the Egyptian president Mubarak at the front of a group of world leaders, where in

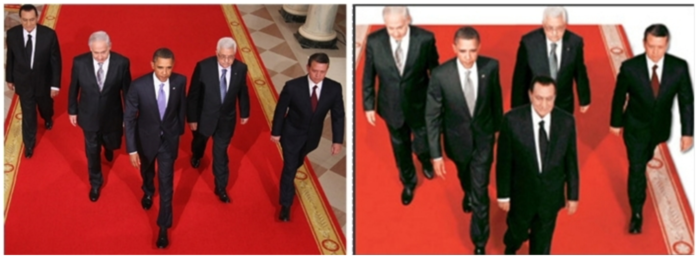

Fig. 1. Image Fakery Examples

the original image he was lagging behind (cf. Figure 1). The fake thus tried to transport a subtle propagandistic message of a distorted reality. In turn, news papers and TV stations had to issue errata to recover their reputation.

Such events are eroding the public trust in media. Therefore, media agencies are required to make their distribution channels capable of identifying media fakery at the earliest stage possible not only to avoid reports of a distorted reality with possible negative consequences, but also to avoid additional costs due to the following correction means. The most desirable solution is automatic fake detection, but current methods still cannot identify semantic inconsistencies in media (cf. [29]). Thus, complementary Web 2.0 community-based approaches were developed to involve people in such processes. Systems such as NewsTrust (http://newstrust.com) pursue such an approach. However, information still has to be pulled by participants, although the current trend hints to real-time requirements and synchronous server side pushes [20] creating a new level of community awareness. Furthermore, the quality of authenticity judgements depends on the trustability of its judges. Current systems establish the trust level of a user by ratings of others which are often subjective and not based on objectively valuable contributions. Furthermore, the willingness to spend time on rating others is mostly not given. Instead of basing trust on subjective opinions, a method is required that objectively adapts trust levels depending on actions.

In this paper we overcome the above problems with an open standard-based collaborative image fake detection system distributed across various communities. The system operates in near real-time and complements traditional automatic approaches. Our approach is powered by a set of Web services based on the MPEG-7 standard as well as by services and infrastructure provided by the open standard Extensible Messaging and Presence Protocol (XMPP) [25, 26] and its extension protocols, in particular XMPP PubSub [17]. A media fake detection application connecting to our infrastructure is realized as a Web 2.0 application consisting of a set of OpenSocial Gadgets [19] for direct communication and the distribution of MPEG-7 [14] multimedia metadata across an XMPP network of media agents.

In Section 2 we first analyze the state-of-the-art of image fakery detection systems and technologies related to our approach. Then, in Section 3 we describe a use case scenario where three media agencies have to detect a faked image, thereby identifying requirements for our system. In Section 4 we present the backend of our system as a multimedia distribution network including its individual parts in detail. In Section 5 we present a media fake detection application powered by our network. In Section 6 we conclude and provide an outlook to further work.

2 Related Work

Faked Image Detection: Faked image detection has been investigated for years and addressed by a number of approaches. Watermarking approaches [24] are based on imperceptibly embedding information within the image content. The requirements of embedding such information in digital images are specially equipped digital cameras. In addition, watermarking degrades the quality of the image content.
In contrast to watermarking approaches, researchers in the field of digital image forensics have developed passive techniques which operate in the absence of any watermark or signature for image authentication (e.g. [6, 21, 11, 31]). They work on the assumption that although digital forgeries may leave no visual clues of having been tampered with, they may alter the underlying statistics of an image that can be detected using statistical models. The major drawback of such tools is that their use in public domains is computationally impractical.
Content based approaches (e.g.[18, 12]) aim at detecting all faked images produced from the original through active manipulation. They are based on similarity search and embed no additional information within the image content, thus considering the image itself as the watermark. The efficiency of such techniques is largely affected by the size of the reference image dataset [18]. Furthermore, current approaches lack discriminative power for fake detection due to the inability of capturing semantic aspects. Our collaborative fake detection system utilizes community aspects in addition to automatic content-based image similarity search techniques [4].
Collaborative Fake Detection: Sharing knowledge and control is the key idea of collaborative fake detection [22]. A Community of Practice [32] is the context where such collaborative activities can be achieved. Knowledge about media is exchanged within the communities of practice for example by the distribution of MPEG-7 metadata [14, 28]. Collaborative judgments and evidence against the suspected fake support the evaluation of semantic inconsistencies that cannot yet be detected with automatic approaches. The important problem faced in collaborative fake media detection is the assessment of trustable authenticity judgments that we address in the scope of this paper.
Trust Management: Trust management is a key issue in distributed networks, especially in sharing environments. Trust provides us with information about the people we should share content with and accept content from. There are some efforts to formalize trust. Massa et al. propose a trust-aware model in which the web of trust is explicitly expressed [16]. Golbeck analyzed and modeled the core characteristics of trust in collaborative social networks and developed several algorithms for computing trust on the example of the TrustMail application

[9]. In this work, we take into account the trust of information sources and the quality of their contributions using a simplified trust mechanism and present a modular trust-aware multimedia distribution network.
MPEG-7: MPEG-7 is a standard for the description of multimedia content. It provides descriptors for various data types - text, graphics, audio, video. In order to achieve interoperability and keep advantages of server side computation we have presented the Lightweight Application Server (LAS) [30] for MPEG-7 Web services. It provides communities with a set of core services and MPEG-7 semantic multimedia metadata and content processing services to connect to heterogeneous data sources [23]. In particular, the LIRE [13] library is used for automatic extraction and indexing of low-level features as well as content based image retrieval (CBIR).
Real-time federation: Due to frequent complaints about the intransparency and lack of control of private data storage with social networking platforms, there are already new alternative platforms emerging (e.g. Diaspora [10]), where the same functionality is offered in a way that anybody can run his instance in federation with others. At the same time, the demand for real-time application behavior [20] speeds up the information flow tremendously. Concepts such as security, privacy and trust have to be weaved in as unobtrusive, transparent, and least blocking as possible. In our approach we aimed to realize these requirements with a network of federation-enabled XMPP servers including respective services and data.
Publish/Subscribe: Nowadays, the *Publish/Subscribe (PubSub)*[2, 5] pattern is omnipresent (e.g. newspapers, blogs, even email lists). There is a *channel of communication* (resp. a *node*), *subscribers* receiving data sent on that channel, and *publishers* who send data *payloads* across the channel. The pattern was also described by Gamma et al. as the behavioural *Observer* pattern [7]. Until today, the pattern is applied successfully, sometimes working locally on one machine or remotely across whole networks. The *XMPP PubSub Extension Protocol*[17] supports the construction of remote PubSub systems transporting XML-based payloads. For this work we demonstrate the distribution of MPEG-7 multimedia content descriptions along with authenticity ratings.

3 Use Case Scenario & Requirements Analysis

In this section we first describe a scenario to understand a media fake detection process in a media distribution network such as in Figure 2. Afterwards, we derive a set of requirements for our system improving the process. Consider the following scenario. A government press agency sends a doctored picture of a successful long-range missile launch to Thompson Reuters as a demonstration of the country's military power, although the real outcome of the event was a crash of the missile. Despite the good cooperation with the government press agency in the past, the responsible media agent recognizes the image content as highly sensitive and thus decides to request expertise on its authenticity before further distribution. Although some trusted experts reviewed the image, the forgery is not discovered, and the picture distributed to customers. TV stations

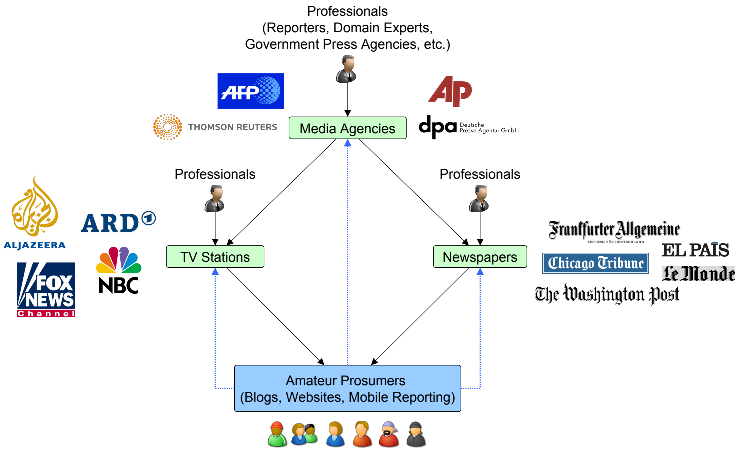

Fig. 2. Examplary excerpt of a media distribution network

and newspapers around the world broadcast the sensitive information to their audiences. After the worldwide publication of the faked picture, a group of local dissidents who eye-witnessed the failed missile launch feel the urge to reveal the truth. In a message sent to Reuters they describe the real situation, send their own picture of the missile crash as proof for their statement, and state their willingness to help prevent such incidents in future. Further expert analysis on both pictures then reveals the fake. As a result, Reuters and all its customers issue a corrective statement to recover their public credibility. However, to prevent further occurrences of such situations, media agents decide to be more cautious towards their information sources or even decide for alternative sources. On the other hand, Reuters acknowledges the group of dissidents' help in discovering the fake and decides to involve their expertise for further authenticity judgement.
From the above scenario we now derive a set of requirements to an information system supporting the process described above, before we explain our approach in the next sections.

– *media & metadata repository*: The first step is to make media and their metadata available for other parties. We base this work on our LAS MPEG-7 services and its repository [30].
– *federated multimedia distribution network*: The most important use case in the scenario is the transport of media (metadata) between entities in real-time. Here, PubSub is the main communication pattern. For a distributed approach, PubSub support is required in a remote and federated manner. The network should support arbitrary payload formats in order to stay generic. Here, we base our approach on the XMPP Protocol and its PubSub extension [17] fulfilling all these requirements.
– *authenticity rating service*: a service is required that allows the collaborative assignment of authenticity ratings to media as well as the computation

and rendering of reasonable aggregates to create awareness for fakes and to support the decision of a media agency to publish a medium or not.
- *trust management service*: a service is required that manages trust relationships between entities again in a federated way and supports the dynamic evolution of trust. Since the service itself must be trusted by its users, privacy and security are non-functional requirements to be guaranteed.

4 A Trust-aware Multimedia Distribution Network

In this section we present a modular trust-aware multimedia distribution network based on the above requirements. In Section 4.1 we describe a basic network building block and its workflow. Each building block implies a simple trust protocol which is formalized in Section 4.2. Finally, we demonstrate the composition of complete information distribution networks of building blocks in Section 4.3.

4.1 The Basic Building Block

Conceptually, the basic building block of our architecture is a variation of the PubSub pattern (cf. Fig. 3). The central parts of this building block are an *untrusted in node* and a *trusted out node* with configuration under control of a *mediator*. For the in node, all of the mediator's *sources* are publishers and subscribers at the same time to support media distribution for collaboration. For the out node, only the mediator is allowed to publish. The list of subscribers reflects the mediator's consumers relying on the authenticity of the information published. First, a source introduces a new medium along with an authenticity

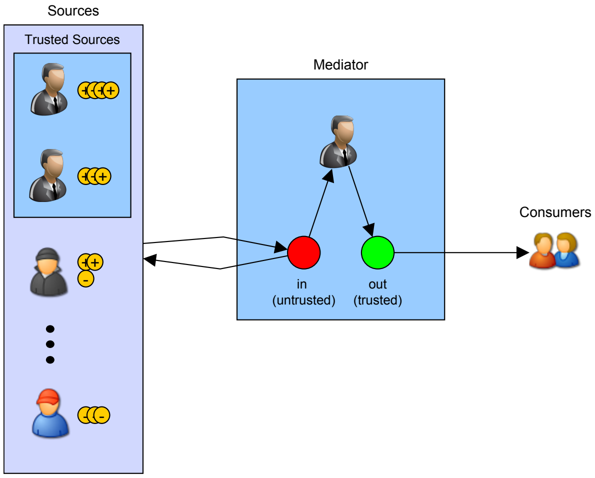

Fig. 3. Building block for media distribution network

rating by publishing it to the untrusted in node that immediately pushes it to all other sources, which in turn publish their authenticity ratings to the same node. Based on ratings from various sources accumulated over time, the mediator eventually decides the information to be trustworthy of being published to the out node or not. The decision depends on the individual levels of trust towards his sources. In Section 4.2 we provide a formalized description of our trust mechanism. Technically, each of the building blocks described above can be realized with a set of components depicted in Figure 4. Any XMPP Server hosting a PubSub service as specified in [17] realizes all necessary functionality regarding the management and configuration of nodes, in particular controlling node subscriber and publisher lists, as well as pushing arbitrary XML-based payloads to subscribers.

4.2 Authenticity Rating & Trust

In this section we formalize the relationship between authenticity ratings and trust used in our approach. Let $J = \{j_1, ..., j_n\}$ be a set of unique identifiers for the entities involved in the fake detection process. For our approach we use JIDs (cf. [25]). Our basic notion of trust involves two entities, i.e. a *trustor* $tr \in J$, a *trustee* $te \in J$ and a *level of trust* $t(tr, te)$ between them. Although there exists work on sophisticated models such as [15], trust-aware social networks usually let users assign a single numerical rating for usability reasons [8]. In our model, the mediator m of a building block from Section 4.1 takes the role of the trustor of a *set of sources* $S_m \subset J$ as its trustees, that publish information payloads i of a certain domain I (in our case the domain of MPEG-7 descriptors).

For authenticity ratings, we define a function r that for a given i and a source s assigns a rating $\in R = \{true, fake\}$. In the following we describe the relationship between authenticity ratings and the dynamic adaptation of trust between involved entities.

Not only is $t(tr, te)$ depending on previous authenticity statements, but also should be adapted dynamically, either reinforcing desirable actions - in our case publishing a faked medium as fake resp. a real one as real - or punishing undesirable actions - in our case publishing a faked medium as real resp. a real one as fake. Thus, reinforcement consists in tr raising his trust level towards te, punishment in lowering it. Thus, each m must be enabled to update trust levels $\forall s \in S_m$. Listing 1.1 sketches an algorithm for updating trust values.

```
trust_update($m \in J, i \in I, x$ action){
  for each $s \in S_m$ {
    if $r(i, s) = fake \wedge x = p_{fake}(m, i)$ then $t(m, s)$++;
    else if $r(i, s) = true \wedge x = p_{fake}(m, i)$ then $t(m, s)$--;
    else if $r(i, s) = fake \wedge x = p_{real}(m, i)$ then $t(m, s)$--;
    else if $r(i, s) = true \wedge x = p_{real}(m, i)$ then $t(m, s)$++;
  }
}
```

Listing 1.1. Updating trust values after publication to trusted out node

Any trust update takes place whenever m feels confident to *publish* i as either *fake* ($p_{fake}(m,i)$) or *real* ($p_{real}(m,i)$). Furthermore, there is the option of *rejecting* any publication on the trusted out node ($rej(m,i)$). In this case, no trust update takes place. Since m in his role as trustor is interested in high-quality media (metadata) and reliable authenticity ratings, he can expose trust levels as an incentive to perform desirable actions only. To decide publication of an i, m relies on ratings from different $s \in S_m$, while using $t(m,s)$ as weighting factor. For a given $i \in I$, a function a returns an aggregate supporting m in his decision which action to take. For simplicity we chose $a(m,i)$ as weighted mean over all ratings on i by $s \in S_m$, where the weights are given by $t(m,s)$ (cf. Equation 1). The intuition behind choosing the weighted mean is that the higher a source's trust value is the more influence his rating has on the resulting aggregate used by m to decide on publication.

$$a(m,i) = \frac{\sum_{j=1}^{|S_m|} t(m,s_j) * r(i,s_j)}{\sum_{j=1}^{|S_m|} t(m,s_j)} \qquad (1)$$

Technically, the dynamic management of trust is realized as a service that maintains individual levels of trust between trustors and their trustees. Ratings of different sources for given information items are covered by another service.

4.3 Construction of a Network

A complete distribution network can now be modeled by reasonably connecting multiple building blocks. The intuition is that each mediator can act as a source for another mediator. Thus, information distribution networks can dynamically evolve over time by simple interactions with XMPP PubSub nodes. It should be noted that it is not necessary that each entity in the network maintains its own XMPP server, which would be acceptable e.g. for a high-profile media agency, but inacceptable e.g. for a freelancing information agent. For these purposes

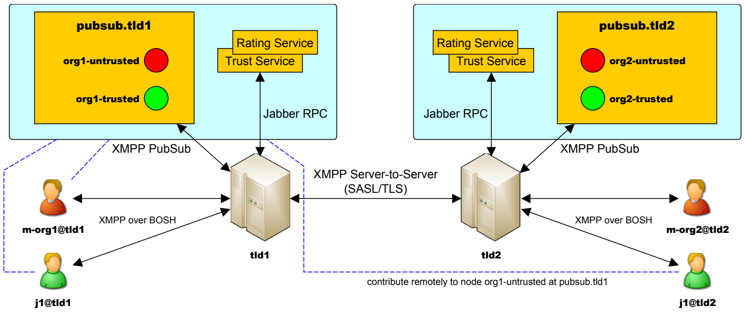

Fig. 4. A Trust-aware Federated Media Distribution Network

it is possible to offer a building block from Section 4.1 as a service, which is hosted on one XMPP server or a whole cluster. On the technical level we realize a network of different interconnected building blocks by a network of XMPP servers in combination with the provision of services for the management of users, communities, MPEG-7 multimedia metadata, trust and authenticity rating as indicated in Figure 4. Given the inherent XMPP server-to-server communication [25, 26], all components are federated and accessible across the network via the protocol and its extensions [1, 17, 27]. In particular, [1] can be used to invoke services of our LAS.

5 A Fake Multimedia Detection Application

In this section, we briefly describe how to apply our trust-aware media distribution network from Section 4 for realizing a fake media detection application. Figure 5 shows a first mockup of such an application consisting of a set of three widgets. In the following we will briefly explain the interface for collaborative fake media detection for both the mediator and his sources, which reflects the workflow from Section 4. In the *Incoming Media Overview*, the user gets an overview of media currently discussed on all untrusted in nodes he is subscribed to. Each element of the list provides a short summary of the medium and its metadata (cf. $i \in I$, Section 4.2) and the weighted authenticity ratings aggregate (cf. function

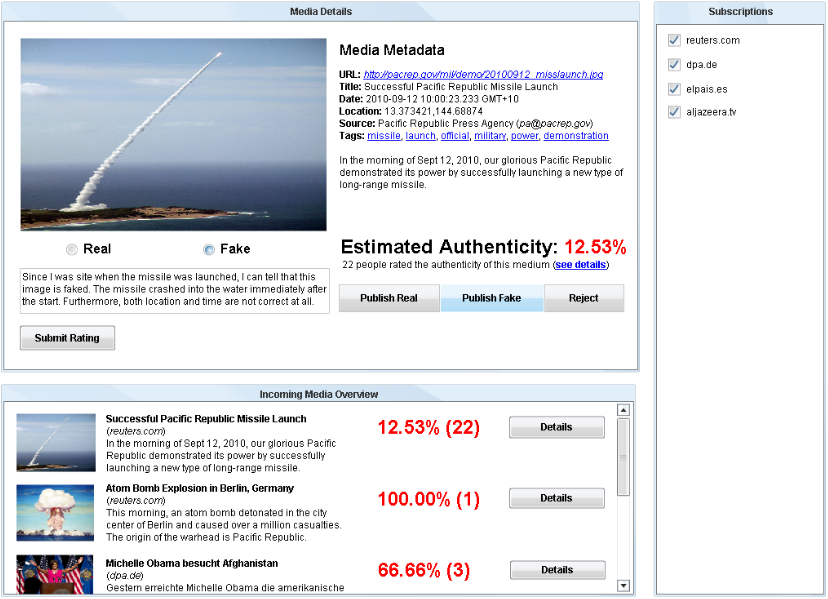

URL: http://pacrep.gov/mil/demo/20100912_misslaunch.jpg
Title: Successful Pacific Republic Missile Launch
Date: 2010-09-12 10:00:23.233 GMT+10
Location: 13.373421,144.68874
Source: Pacific Republic Press Agency (*pa@pacrep.gov*)
Tags: missile, launch, official, military, power, demonstration

In the morning of Sept 12, 2010, our glorious Pacific Republic demonstrated its power by successfully launching a new type of long-range missile.

Since I was site when the missile was launched, I can tell that this image is faked. The missile crashed into the water immediately after the start. Furthermore, both location and time are not correct at all.

Estimated Authenticity: 12.53%
22 people rated the authenticity of this medium (**see details**)

Successful Pacific Republic Missile Launch
(*reuters.com*)
In the morning of Sept 12, 2010, our glorious Pacific Republic demonstrated its power by successfully launching a new type of long-range missile.
12.53% (22)

Atom Bomb Explosion in Berlin, Germany
(*reuters.com*)
This morning, an atom bomb detonated in the city center of Berlin and caused over a million casualties. The origin of the warhead is Pacific Republic.
100.00% (1)

Michelle Obama besucht Afghanistan
(*dpa.de*)
Gestern erreichte Michelle Obama die amerikanische
66.66% (3)

Fig. 5. Widget-based UI of a Multimedia Fake Detection Application

a, Section 4.2). From this list, the user can select any element, which is then rendered in more detail in the *Media Details* widget. Apart from the medium and its metadata, the user finds different buttons, depending on his role. As an information source, the user finds a rating interface, which allows him to choose between real or fake, add a comment and submit his rating. On submission, a triple consisting of a source identifier, a media identifier and a rating is encoded as an XML payload and published to the in node again. After automatic forwarding to all subscribers, their interfaces are updated with the new information. As a mediator, the user can decide on the three different actions p_{real}, p_{fake}, and rej (cf. Section 4.2) by pressing the respective buttons. A trust update (cf. Listing 1.1) is executed after any publication to the trusted out node by invoking the respective LAS service. Due to space restrictions, we will not elaborate here on further UI elements, such as media annotation (cf. [23]), advanced trust visualisation, etc.
Technically, the interface is realizable as a set of OpenSocial [19] gadgets using XMPP/LAS AJAX client libraries to connect to the XMPP server network and its services. For the access to PubSub nodes, we implemented an extension of the dojo XMPP library realizing the most important use cases of [17]. For the access to LAS Services, we implemented an AJAX connector client library. However, a further extension of the dojo XMPP library realizing the Jabber RPC extension protocol is a preferable alternative for the future.

6 Conclusions

In this paper we have demonstrated an approach for collaborative fake media detection based on a federated, trust-aware media distribution network with near real-time properties. We have presented an overview of related work in the domain of fake media detection, which is dominated by image processing approaches, that still do not bridge the semantic gap [29] and by community approaches lacking real-time communication and trust adaptations based on objective actions. Thus, we proposed our approach to overcome these challenges. Starting from a realistic use case scenario we elicited requirements and presented a realization as an XMPP-based and Web service-enhanced multimedia distribution network supporting arbitrary XML-based payload format. Finally, we sketched the design of a Web-based fake media detection application taking benefit from our network and its services.

At the time of writing this document many components of our multimedia distribution network as well as connector clients have been realized and evaluated. We already gained experience with XMPP-enabled OpenSocial Gadgets and therefore extended the well-known dojo JS library with support for PubSub, multi-user chats, etc. [33]. With these extensions, a real-time microblogging application was easily realizable. Although the XMPP standard provides detailed documentation about the protocol itself, there is not too much information which PubSub node topologies are suitable when scaling up to larger and highly distributed networks. Thus, we are currently evaluating architecture scalability and performance in the context of the ROLE project (http://role-project.eu), where XMPP also serves as an open standard infrastructure for Widget-based PLE (Personal Learning Environments). Currently, we realize the fake multimedia detection application based on the design presented in the context of this work.

Acknowledgments. The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no 231396 (ROLE project) as well as DAAD.

References

1. D. Adams. XEP-0009: Jabber-RPC. Technical report, XMPP Standards Foundation, February 2009.
2. K. P. Birman and T. A. Joseph. Exploiting Virtual Synchrony in Distributed Systems. In *SOSP '87: Proceedings of the eleventh ACM Symposium on Operating systems principles*, pages 123–138, New York, NY, USA, 1987. ACM.
3. H. Chesbrough. *Open Innovation: The new imperative for creating and profiting from technology.* Harvard Business School Press, 2003.
4. R. Datta, D. Joshi, J. Li, and J. Z. Wang. Image retrieval: Ideas, influences, and trends of the new age. *ACM Comput. Surv.*, 40:5:1–5:60, May 2008.
5. P. T. Eugster, P. A. Felber, R. Guerraroui, and A.-M. Kermarrec. The Many Faces of Publish/Subscribe. *ACM Computing Surveys*, 35(2):114–131, June 2003.
6. J. Fridrich, D. Soukal, and J. Lukáš. Detection of Copy-Move Forgery in Digital Images. In *Proc. of DFRWS 2003*, pages 90–105, 2003.
7. E. Gamma, R. Helm, R. E. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software.* Addison-Wesley, Reading, MA, 1995.
8. J. Golbeck. Trust and Nuanced Profile Similarity in Online Social Networks. *ACM Transactions on the Web.* In press.
9. J. Golbeck. *Computing and Applying Trust in Web-based Social Networks.* PhD thesis, University of Maryland, 2005.
10. D. Grippi, M. Salzberg, R. Sofaer, and I. Zhitomirskiy. diaspora*. The privacy aware, personally controlled, do-it-all, open source social network, 2010.
11. M. Johnson and H. Farid. Exposing Digital Forgeries by Detecting Inconsistencies in Lighting. In *ACM Multimedia and Security Workshop*, pages 1–10, New York, NY, USA, 2005. ACM.
12. Y. Ke, R. Sukthankar, L. Huston, Y. Ke, and R. Sukthankar. Efficient Near-duplicate Detection and Sub-image Retrieval. In *MM '04: Proceedings of the ACM international conference on Multimedia*, pages 869–876, NY, USA, 2004. ACM.
13. M. Lux and S. A. Chatzichristofis. Lire: lucene image retrieval: an extensible Java CBIR library. In *MM '08: Proceedings of the 16th ACM international conference on Multimedia*, pages 1085–1088, New York, NY, USA, 2008. ACM.
14. B. S. Manjunath, P. Salembier, and T. Sikora. *Introduction to MPEG-7, Multimedia Content Description Interface.* John Wiley and Sons, Ltd., June 2002.
15. S. Marsh. *Formalising Trust as a Computational Concept.* PhD thesis, University of Stirling, Department of Mathematics and Computer Science, 1994.
16. P. Massa and B. Bhattacharjee. Using Trust in Recommender Systems: An Experimental Analysis. In *Proceedings of iTrust2004 International Conference*, pages 221–235, 2004.

17. P. Millard, P. Saint-Andre, and R. Meijer. XEP-0060: Publish-Subscribe. Technical report, XMPP Standards Foundation, July 2010. Draft Standard.
18. S. Nikolopoulos, S. Zafeiriou, and N. Nikolaidis. Image replica detection system utilizing R-trees and linear discriminant analysis. *Pattern Recognition*, 43(3):636–649, March 2010.
19. OpenSocial Specification 1.0. Technical report, OpenSocial and Gadgets Specification Group, March 2010.
20. T. O'Reilly and J. Batelle. Web Squared: Web 2.0 Five Years On, 2009. Special Report Web 2.0 Summit.
21. A. Popescu and H. Farid. Statistical Tools for Digital Forensics. In *Proceedings of the 6th International Workshop on Information Hiding*, pages 128–147, Toronto, Canada, 2004. Springer-Verlag, Berlin-Heidelberg.
22. K. A. N. Rashed and R. Klamma. Towards Detecting Faked Images. In A. Carreras, J. Delgado, X. M. as, and V. Rodriguez, editors, *Proceedings of the 11th International Workshop on Interoperable Social Multimedia Applications(WISMA10), Barcelona, Spain, May 19-20*, May 2010. CEUR Workshop Proceedings, Vol. 583.
23. D. Renzel, R. Klamma, Y. Cao, and D. Kovachev. Virtual Campfire - Collaborative Multimedia Semantization with Mobile Social Software. In R. Klamma, H. Kosch, M. Lux, and F. Stegmaier, editors, *Proceedings of the 10th International Workshop on Semantic Multimedia Database Technologies (SeMuDaTe'09), Graz, Austria*, 12 2009. CEUR Workshop Proceedings, Vol. 539.
24. C. Rey and J.-L. Dugelay. A survey of Watermarking Algorithms for Image Authentication. *EURASIP J. Appl. Signal Process.*, 2002(1):613–621, 2002.
25. P. Saint-Andre. RFC 3920 – Extensible Messaging and Presence Protocol (XMPP): Core. Technical report, Jabber Software Foundation, October 2004.
26. P. Saint-Andre. RFC 3921 – Extensible Messaging and Presence Protocol (XMPP): Instant Messaging and Presence. Technical report, Jabber Software Foundation, October 2004.
27. P. Saint-Andre. XEP-0045: Multi-User Chat. Technical report, XMPP Standards Foundation, July 2008. Draft Standard.
28. G. Shih-Fu Chang, A. B. Chan, and P. J. Moreno. Overview of the MPEG-7 standard. *IEEE Trans. Circuits and Systems for Video Technology.*, 11(0):688 – 695, 2001.
29. A. Smeulders, M. Worring, S. Santini, A. G. A, and R. Jain. Content-Based Image Retrieval at the End of the Early Years. *IEEE Trans Pattern Anal Mach Intell*, 22(12):1349 – 1380, 2000.
30. M. Spaniol, R. Klamma, H. Janßen, and D. Renzel. LAS: A Lightweight Application Server for MPEG-7 Services in Community Engines. In K. Tochtermann and H. Maurer, editors, *Proceedings of I-KNOW '06, 6th International Conference on Knowledge Management, Graz, Austria, September 6–8,2006*, J.UCS (Journal of Universal Computer Science)Proceedings, pages 592–599. Springer-Verlag, 2006.
31. W. Wang and H. Farid. Exposing Digital Forgeries in Video by Detecting Double Quantization. In *ACM Multimedia and Security Workshop*, pages 39–48, Princeton, NJ, September 2009. ACM.
32. E. Wenger. *Communities of Practice: Learning, Meaning, and Identity.* Cambridge University Press, Cambridge, UK, 1998.
33. M. Wolpers, M. Friedrich, R. Shen, C. Ullrich, R. Klamma, and D. Renzel. Early Experiences with Responsive Open Learning Environments. In H. Maurer, N. Kulathuramaiyer, and K. Tochtermann, editors, *Proceedings of I-KNOW 2010, 1-3 September 2010, Graz, Austria*, pages 391–402, 2010.