

---

## Decentralised social network management

---

Stefano Bortoli\*, Themis Palpanas and  
Paolo Bouquet

Department of Information Engineering and Computer Science,  
University of Trento,

Via Sommarive 14, 38100 Povo, Italy

E-mail: [bortoli@disi.unitn.it](mailto:bortoli@disi.unitn.it)

E-mail: [themis@disi.unitn.it](mailto:themis@disi.unitn.it)

E-mail: [bouquet@disi.unitn.it](mailto:bouquet@disi.unitn.it)

\*Corresponding author

**Abstract:** Social networking sites have gained much popularity in the recent years because of the opportunities they give to people to connect to each other in an easy and timely manner, and to exchange and share various kinds of information. However, these sites are based on a centralised paradigm, which limits the mobility of their users, and ultimately, their chances to establish new relationships and benefit from diverse networking services. In this paper, we argue for a decentralised paradigm for social networking, in which users retain control of their profiles, and social networking sites focus on the delivery of innovative and competitive services. In this environment, both the social networking sites and their users will be able to develop to their full potential. This goal can be achieved by using a combination of semantic web technologies and tools, loosening the bind between the social network management and social networking web applications.

**Keywords:** unique identification; web application; entity name system; FOAF; decentralised social network management; web-based communities; social network portability.

**Reference** to this paper should be made as follows: Bortoli, S., Palpanas, T. and Bouquet, P. (2011) 'Decentralised social network management', *Int. J. Web Based Communities*, Vol. 7, No. 3, pp.276–297.

**Biographical notes:** Stefano Bortoli is a PhD candidate at the ICT International Doctoral School of the University of Trento. He received his MSc with honour at the University of Trento (2006) and MSC at the Free University of Amsterdam (2007). He worked as a Research Assistant as part of the Data and Knowledge Management Group of the Department of Information Engineering and Computer Science of the University of Trento since 2004. In particular, he was involved in the EU founded projects VIKEF and OKKAM. In this period, he actively participated in research activities, the results of which have been presented in national and international conferences and workshops.

Themis Palpanas is a Faculty Member of Computer Science at the University of Trento, Italy, working on data management and data analysis. He received his BS from NTUA, Greece, and PhD from the University of Toronto, Canada. Before joining the University of Trento, he worked at the IBM T.J. Watson Research Center. He has also worked for the University of California, Riverside, and visited Microsoft Research and the IBM Almaden Research Center. His research solutions have been implemented in world-leading

commercial data management products and he is the author of five US patents. He is a Founding Member of the Event Processing Technical Society and is serving on the Editorial Advisory Board of the *Information Systems Journal* and as an Associate Editor in the *Journal of Intelligent Data Analysis*. He has served on the programme committees of several top database and data mining conferences, and is the General Co-chair for VLDB 2013.

Paolo Bouquet is an Associate Professor in Computer Science at the University of Trento since 2004. He has a Master in Philosophy (University of Milan) and a PhD in Philosophy of Science (University of Genoa). He was a visitor at Stanford University (1994) and at the University of Texas at Austin (1995). In 1998, he was awarded a Visiting Research Fellowship of the Caledonian Research Foundation – Royal Society of Edinburgh (Scotland) to carry research at the University of Dundee. His main research areas are: knowledge representation and management, semantic interoperability, formal theories of context and contextual reasoning. His results are published as articles in several international journals and were presented as papers at many international conferences. He has been responsible for the research unit of the University of Trento of several projects, including: OKKAM (Project Coordinator), VIKEF, WISDOM and EDAMOK. He organised several international workshop and conferences.

---

## 1 Introduction

A social network is a structure composed by individuals or organisations that are connected by one or more specific types of interdependency, such as values, ideas, friendship, etc. In the past years, we have witnessed a fast growth of web-based social network services aiming at building online communities of people willing to share their interests or activities (Freeman, 2006).

During the early phase of social network service development, the focus was on connecting people tied by specific kind of relations, for example, ex-classmates. Recently, some social network services left this bind to specific relation between social-beings offering a more general purpose social networking service. Furthermore, this kind of services allows people to build communities or groups of interest about specific topics, potentially satisfying the social networking needs of every user. This approach led to an impressive growth of few social network services involving hundreds of millions of users often living in localised geographical area. Among others, the most successful web-based social network services are Facebook, Myspace and Orkut.

This fast growth in the number of social network services users created a brand new business market related to commercial space and market analysis. This is one of the reasons that led social network providers to the definition policies aiming at reducing the ‘mobility’ of users through the several social network service applications. Namely, a social network user who finds a new interesting social network platform is discouraged from moving on such a platform because she would have to repeat all the preliminaries operations necessary to initiate their social network activities. Among others, the most time expensive and sometime cumbersome tasks are the description of a personal profile and the ‘reconstruction’ of the social relations binding the user to the other social-beings.

In this context, we now want to move towards a model where individual users and their relations become the focus of social networking applications. In this new model, users have more control over the data/information they own, such as the data that refer and are relevant to them, including the information relative to their social relations. Furthermore, users should be empowered to perform some basic, yet fundamental, action on these data. An example is the capability of moving social network information (i.e., the connection with other social network users) across different social network applications, keeping control over this information.

In this paper, we argue in favour of a decentralised paradigm that allows users to move across web-based social networking applications freely and easily, giving users more power and control over their data. Furthermore, when social network applications enable the complete portability of social networks, several advantages come for social networking service providers.

First, data about users and their connections are under the control of the users, which means that (in principle) they can be uploaded into any social networking platform that is compatible with the format in which these data are structured. This way, users can benefit from the different services offered by a variety of social networking platforms. Second, social networking applications can focus on the task of providing value added services to users, and try to attract new users from a pool that is now much larger. By allowing users to establish relationships across different social network sites that cater to social networks with different interests and goals (e.g., artistic, professional, etc.), all social networking platforms can benefit from the increased mobility of users. Furthermore, if a user is free to transfer social network information across different platforms, the 'extended' social network brings in more users by minimising the *new user* entry cost. Similar observations to the ones we made above have also appeared in other studies (Ramakrishnan and Tomkins, 2007).

In order to achieve these goals we have implemented a tool enabling users to manage personal social network profiles defined using the well known Friend-Of-A-Friend (FOAF) ontology as standard vocabulary (Briceley and Miller, 2004). Any social network profile presents only the FOAF description of a person and her social network (i.e., the list of known persons). Thereby, in order to assure an automatic social network integration with other social network profiles forming an extended social network graph we need to find a way to identify uniformly all the people across different profiles. To accomplish this objective we enabled our tool to interact with an entity name system (ENS) providing shared globally unique identifiers. This way, any user can define and manage her social network in a social network web application independent environment, using a standard vocabulary, and with the possibility of a smooth integration with other social networks based on person globally unique identification. Such social network profiles can in principle be used to access any social network application, making possible concrete social network portability. In order to complete the picture outlined so far, we consider the existence of some kind of web accessible repository where the user can upload and manage social network profiles. This repository would have to provide access control functionality, preventing unwanted social network information disclosure. Thus, a user that wants to access a web-based social networking application should simply provide the link to the social network profile repository and provide the credentials enabling the application to read/write the profile.

The remaining of this paper is organised as follows. We first describe a scenario motivating our work in the context of social network portability and decentralised management in Section 2. In Section 3 we provide an analysis of other approaches aiming at social network portability. In Section 4, we describe the foundations of the proposed paradigm, then described in Section 5. In the latter we also present a neutral web-based tool supporting users in editing social networks and integrating globally unique identifiers. Finally, we conclude in Section 6.

## 2 A motivating scenario

In order to make explicit our vision about the future of web-based social networking and related communities, we describe in the following a motivating scenario. Starting from this scenario we also argue about how a social network user can manage social network activities across different web-based social networking applications according to the nature of the relation binding her to other social beings.

The intuition behind our approach can be explained through a simple analogy:

“Imagine you go to a pub, and meet new friends. Even though the reason for meeting in that specific pub can be to drink a particular brand of beer, you and your friends are still the ‘owners’ of your relationship, and such a relationship should be independent of the fact that you met in that pub. For example, one day you may decide to watch a football match with the same friends. Then you and your friends can simply meet at another pub that screens this particular football match. Hence, the two pubs stay the owners of the services they offer (serving good beer, or having a large screen for football matches), and they compete on the basis of the services they offer. Nevertheless, pub customers are not locked-in, which means that they can decide to use other services (or even the same services, but of better quality) in another pub.”

In a similar way, we foresee a system where a user can access a social network application offering a specific kind of service, meet new friends, and be able to move this new relation from the original social network application to others.

Imagine now a system where a user can manage her social network information in a neutral environment, outside any of the existing social network application. In this context, the user can edit personal information, add, remove and edit different kind of relationships (e.g., friendship), and guarantee access to this information only to trusted applications. Every social network editing operation is automatically reflected into the social network web application subscribed by the user, and, in turn, every editing operation performed by the user through the social network application can be reflected in the neutral environment. Notice that in principle the user can manage different social networks according to the context, and subscribed social network application are enabled to access information only if explicitly authorised.

Assume that the user subscribed to a social network application for online poker players. This application allows the user to build small communities of players, and organise matches as well as tournaments. The social network profile describing the user social network is safely stored in an authorised repository. Furthermore, imagine that the user then heard about a new social network application supporting interesting, novel features for poker players. The user subscribes to the new social network, which enables the application to access the poker-related social network information stored in the

repository. Any poker-mates previously subscribed to the new social network application, are automatically added to the account of the user. The rest of the poker-mates known in the old social network application are invited to subscribe the new one, and with very little effort the user can play poker with the old group of poker-mates in the new environment. New poker players may have also joined the small community, extending the poker-related network of the user. At a later time, the first social network application may attract back our user. The user can effortlessly restart using it, along with the new friends met in the other social network application.

In order to make the above scenario real, there are several prerequisites to be considered:

- *Use of shared ontology.* All the social network profiles must be structured using ontologies describing the information within the profiles, as well as the relationships among them. The best option is that users share the same ontology to describe their social network profiles (even though techniques for aligning different ontologies are also available) (Euzenat and Shvaiko, 2007). Using shared ontologies allows different systems to share a common understanding of the information contained in a social network profile.
- *Uniform identification system.* Every social network user must be identified uniformly across the various social network profiles. This fact would enable automated and smooth social network integration. If two social network profiles present a description of a person, and this person is identified in the same way in both profiles, then it is possible to correctly integrate the two social networks and trigger the proper action.
- *Existence of neutral toolkit.* The largest part of social network information is usually managed within the social network application boundaries. It is obvious that we need to enable users to manage their social network also outside this ‘walled garden’. This toolkit should include decentralised social network repositories for storing the required information.
- *Definition of fine-grained access control policies.* The user must be able to select which applications can access and edit her social network information. Towards this goal, we need an access control protocol comparable to the one used for e-banking. This way, the social network user, who is the owner of the social network information, can give permission to the social network web application to access only the information necessary for social network activities in a safe and controlled environment.

### 3 Related work

Social network portability is an open discussion ongoing into the WWW community. In this section we present an overview of the known approaches aiming at making social networks distributed and portable.

Tim O’Reilly introduces the concept of ‘social network operating system’ in (O’Reilly, 2007). In this article the author foresees a system enabling the portability of social networks and personal information from social network application to others. O’Reilly points the uselessness of repeating the personal and social network description

process any time a user access a new social network application. He aims at the definition of a set of user-centric application enabling the description of personal social network, focusing then on definition of tools enabling the automatic integration of the described information with other sources. This article presents mostly a vision and does not present any concrete solution.

Tungare (2007) outlines a decentralised paradigm for social network web application interaction that uses a RESTFull protocol enabling a bi-directional communication between social network applications. The motivation and advantages of the system sketched out are very similar to what we propose in this study, but what Tungare proposes does not deal with social network representation and identifications issues.

Kistner (2007) describes a possible architecture for decentralised social networks. The proposed system requires an identity layer composed by the following elements: identifier database, identifier data, and authentication protocol. The identification layer would provide the basic features to implement a decentralised social network based on four pillars: contact, contract, content, and context (4Cs). The interaction between social network users relies on a subscription system (contact), enabling the owner of content to decide which operations other users can perform on it (contract). This information can then be enriched with contextual information, or context related services.

In the rest of this section, we discuss in more detail the relevant efforts that are under way in the FOAF community, as well as initiatives on new web data representation formats, and novel platforms and architectures.

### 3.1 FOAF related projects

FOAF is a machine-readable ontology describing persons, their activities and their relations to other people and objects. The FOAF project, which defines and extends the vocabulary of a FOAF profile, can be considered the first open standard for social semantic web application in that it combines RDF technology with 'Social Web' concerns (FOAF Project, 2000). One of the first tools enabling users to create FOAF profiles was 'foaf-a-matic' (<http://www.ldodds.com/foaf/foaf-a-matic>). This application essentially is a form enabling users in describing themselves using FOAF properties, listing known persons, generating in the end RDF FOAF profiles. The main weak point of this application is the fact that does not support the editing of the profile once created. Thus, any modification of the generated profile must past through a manual editing of the FOAF RDF file or a new creation from scratch. Another description of a potential solution for social network portability based on FOAF is presented in (Bojars et al., 2008). In this paper the authors mainly formalise the idea behind the FOAF project, providing a more detailed description about how a social network could be described using the FOAF and SIOC ontologies (Breslin et al., 2005). The paper does not provide any solution for the concrete maintenance of both social network profiles and how these profiles should share the same URI to identify the same person.

A decentralisation of social network system based on FOAF is proposed by Yeung et al. (2008). In this paper, the authors sketch a system where users manage all social network information on a trusted server, relying on some access control system to enable social network application to access and use the 'decentralised information' for social network activity. The central point of the paper is to give to the user the power to manage their data and provide a central access point for social networking. The weak points of this approach are two. First the trusted server information repository must be able to

guarantee access and manage all information needed to realise the social network services. This, in our opinion limits the evolution of social network services as, in order to access new services in one application, the user should set his social network information in a way that the social network application can use it to realise the offered service. Thus, any time a new service is offered, the central social network information repository must be updated in order to enable the access to such a new service. We believe that users should own the social network information, but that social network service provider should keep, respecting privacy issues, all the information necessary to provide specific services. Social network providers should compete on offered services in order to guarantee technological evolution. Second the social network integration relies on a weak identification system as URL pointing at FOAF profiles and inverse functional properties are volatile and do not guarantee complete social network integration (Bortoli et al. 2007).

Sapkota et al. (2005) propose a FOAF-based peer-to-peer social network system. The authors outline a peer-to-peer system where users can handle their social network in personal profiles described according to an extended FOAF ontology. The main advantage of this solution is that the user can manage trust and personal information disclosure in a FOAF-based profile completely under his domain. The main drawback is that the system is conceived as peer-to-peer system to manage social network as community of peers without considering the services built on top of it. Namely, the user can access and manage in a decentralised fashion social networks, but the services layer is missing in the big picture.

### 3.2 *Web data format initiatives*

In the context of the Microformat initiative (Allsopp, 2007), a social network portability based on microformat is under definition. The proposed social network portability system, named XHTML Friends Network (XFN – <http://gmpg.org/xfn/>), relies on hCard + xfn microformats for people and relations description (Andrews, 2003). The social network description passes through HTTP references pointing at web pages (e.g. blog, personal pages etc.) of friends. Such web pages should then in turn embed XHTML hCard (<http://microformats.org/wiki/hcard>) presenting the person description and further reference to friends. It is possible to notice how this approach replicates the idea of FOAF using an XHTML syntax rather than resource description framework (W3C, 1999). Microformats have the general advantage of being more easily understood and handled by WWW marketers and developer. However, microformats in general inherit all the limitations of XML when compared to RDF. Namely, in microformats one can tell that a string is a ‘name’ but one cannot tell whose name it is. RDF provides a triple structure enabling the definition of statement of the type (*GLOBAL\_Id, hasName, John Smith*) enabling an exact disambiguation of which John Smith we are talking about. Furthermore, microformat is by design not formally extensible, and cannot be used to perform automatic reasoning steps where possible. Finally, in order to define a social network on base of the proposed microformats, the promoter of the project assumes that every member of the social network owns and maintains a web page, where they publish social network information. We believe this assumption is not valid for a large part of the social network users.

DataPortability is a collaborative effort promoted by a group of volunteers and WWW marketers to enhance the capability of controlling, sharing, and moving data across independent systems. DataPortability endorses the idea that users should be able to move, share, and control for example their identity and all the relevant personal data as photos, videos, conversations etc. The project aims at defining best practices for integrating existing open standards and protocols. The main goal is to enable end-to-end data portability between online tools and services. This frictionless movement through the network of services should foster stronger relationships between people and services providers and help in building a healthy networked ecosystem (Bizannes et al., 2009).

Open data definition (ODD) initiative developed in the context of the Elgg project (<http://elgg.org/>), offers a framework supporting data portability across different social network applications. Basically, ODD provides a neutral format to represent people, sites, 'objects', and the relations among them. This format can be used to implement an export/import tool enabling the migration of social networks across social networking web applications supporting it. The solution proposed seems valid from the technical point of view. Nevertheless, this approach does not rely on semantic web standards (such as RDF) and deals only marginally with the entity identification problem.

OpenSocial (<http://www.opensocial.org/>) is an initiative aiming at providing a set of common application programming interfaces (APIs) for web-based social network applications. These APIs are developed by Google along with MySpace and a number of other social networks providers. OpenSocial is commonly described as a more open cross-platform alternative to the Facebook platform. The main aim of OpenSocial is to create an ecosystem of social network application sharing common software APIs and templates. The main disadvantage of this kind of system is that it relies on software API, which cannot be considered a general solution for data portability (Google Press, 2007).

### 3.3 Novel platforms and architectures

The architecture of a platform for social network interoperability is described by Mostarda et al. (2008). In this paper the authors present the description of a software architecture enabling interoperability between different social network web applications. The work relies on two main pillars: first, the user has an 'identity page' presenting all the social network information according to the favourite standard, the identity page URL works as identifier for the user; second, OpenId as single sign system to access all the social network applications. The platform would then provide connectors and converters enabling all the social network application to access the information in the user identity page according to the required format. In other words, the authors propose a sort of centralised 'social network information' repository providing proper API supporting social network application with secure access to the user data. This project presents solution to many practical problems, but let unresolved the problem of identification of person within the user social network. The above work does not present any specific solution to the automatic social network integration problem once the user information is imported into the social network application.

Seok-Won, S. et al. (2009) present a technical report about the implementation of platform architecture for decentralised social networking. The authors describe a system where user can retain control over data which are kept in distributed, decentralised, storage supports handled by 'butlers'. Application can then access the distributed data through open APIs which integrates fine grained access verification to avoid unwanted

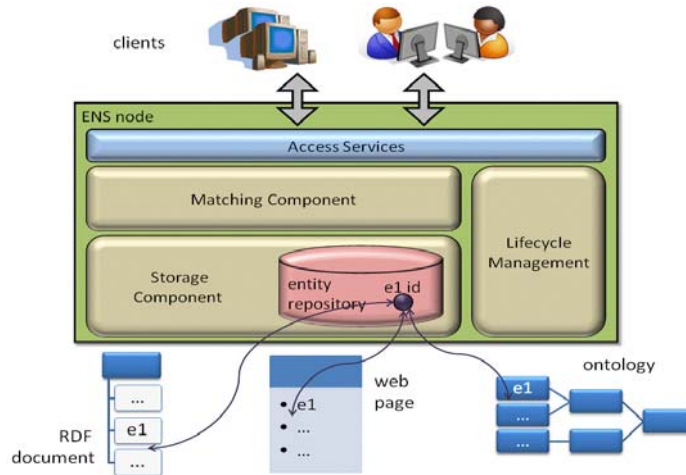


personal information disclosure. The system relies on OpenId as provider for identifier as well as for Certificate Authority to guarantee trusted communication and query propagation across the distributed personal data butlers. The person centric approach adopted by the authors leads to a very promising development of decentralised social networking activities, but presents at least a weak point at the state of the art. Indeed, a user can define personalised relations among data extending without any formalism the ontology used to represent and store the data. This approach on one hand gives lot of freedom to the user, but on the other hand compromises a smooth social network portability and integration when required. Furthermore, this can create problems during query resolution. For example, potentially available information cannot be retrieved because data are described using different names between for properties relating personal data.

#### 4 Preliminaries

We now give a brief overview of the ENS (a more detailed presentation can be found in Bouquet et al., 2008, 2009; Bazzanella et al., 2008), which we will use as the basis for our discussion. Note however, that our discussion is relevant to any system for entity identification management. The overall goal of the ENS is to handle the process of assigning and managing unique identifiers for entities in the WWW. These identifiers are global, with the purpose of consistently identifying a specific entity across system boundaries, regardless of the place in which references to this entity may appear.

**Figure 1** Schematic view of the ENS and its interactions (see online version for colours)



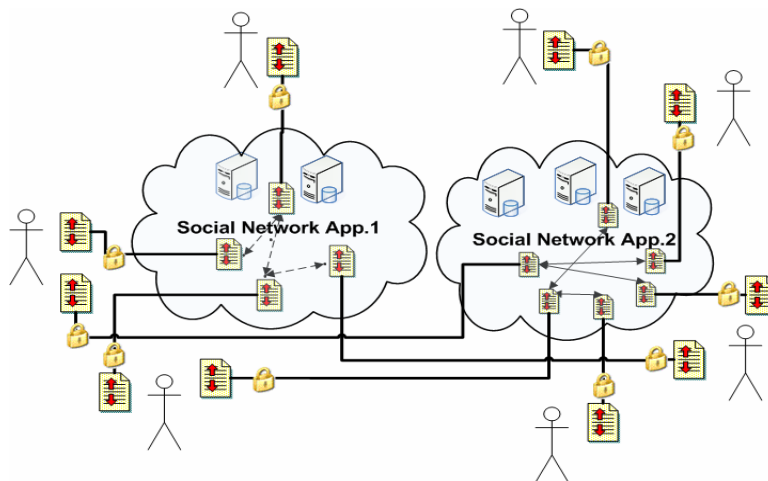
The ENS has a repository for storing entity identifiers (note that this repository will be distributed and replicated) long with some small amount of descriptive information for each entity. The purpose of storing this information is to use it for discriminating among entities, not exhaustively describing them. Entities are described by a number of attribute-value pairs, where the attribute names and the potential values are user-defined (arbitrary) strings. Clients interact with the system through the Access services layer.

Clients can be both human users and applications, and may inquire about the identifier of an entity by providing a set of attributes that describes this entity. If the entity exists in the repository, the system returns its identifier. Clients may also modify the state of the repository, either by inserting a new entity in the system, in which case the ENS returns the newly assigned identifier, or by changing some of the attributes of an existing entity. As shown in Figure 1, the end result is that all instances of the same entity (i.e., mentioned in different systems, ontologies, web pages, etc.) are assigned the same identifier. Therefore, joining these documents and merging their information becomes a much more simple and effective process than before.

## 5 Towards a decentralised paradigm

In this section we present in detail the new paradigm we propose for social networks and their portability. The first point to analyse concerns social network ownership. We believe that a social network (i.e., the set of connection between the user and the tied people) only belongs to the person describing it. When a user defines a new relationship with an unknown person in a social network application environment, the user should be the only owner of this information. Stating the contrary is comparable to constraining the realisation of social activity only to the ‘place’ where the two persons got in contact for the first time. For this reason we think it’s essential to define an open standard allowing the description of social network of a person enabling its complete portability. The social network information should not be visible and understandable exclusively in the application in which the social relations between social-being is expressed. Users should be able to describe in a generic machine-processable way their own social network only once, and in a single document under their complete control.

**Figure 2** An overview of a possible structure of the decentralised social network paradigm (see online version for colours)



Social network information should be decentralised with respect to the social network application, and managed by the user as a private piece of knowledge. Any user should be allowed to manage/edit his/her social profile(s) out of any social network application.

The user 'loyalty' to a social network application should not rely on concerns about losing contact with tied people, but should be based on the quality of the services offered to manage social relations. A view over a possible schema representing this new paradigm is depicted in Figure 2.

In this sense, what the FOAF project proposes seems to be a suitable solution for the definition of decentralised and portable social networks. In fact, what FOAF proposes is essentially to describe social networks in a machine-processable way. Users are asked to define social network personal profile and make it publicly available as any other web document. However, this solution considered in isolation suffers from the following drawbacks.

- *Lack of access control.* Openly sharing social networks data raises privacy concerns related to information access control. Indeed, users might enjoy on the one hand the advantages of social network portability, but on the other hand, they are not able to control which application is going to access their information. In our opinion this fact limits the adoption among average social network users, stressing again the private nature of social network information.
- *Weak identification system.* The use of FOAF ontology (as any other formally defined ontology) to define social network profile guarantees a uniform interpretation of the attribute describing the person and their connection. Unfortunately, FOAF alone suffers of some identification issue. In fact, the identification of persons in FOAF relies on the fact that different people would use the same values for specific kind of properties (i.e. inverse functional properties). It is quite natural that the values of such properties might not be shared uniformly among all the persons who want to refer to the same friend in their FOAF social network profiles. This decreases the concrete capability of automatically connect and integrate social networks (Bortoli et al. 2007).
- *Lack of expressive social network ontology.* The FOAF ontology presents only the 'knows' relation to tie two persons in a social network. This relation has the good property of being general enough to satisfy any kind of social network relation. Nevertheless, the capability of specifying the nature of the relation connecting two persons is quite important. It is actually the base of the definition of social network itself. Any subject is indeed part of different, probably disjoint, social networks defined according to the founding relations connecting him with other subjects. The fact that FOAF presents only one kind of relation prevents users to adopt it to specify different social network according to the context involving them. For example, it's quite common to discern and keep separated work social network from friends' social networks. Nevertheless, FOAF can be considered a good starting point for the definition of more context specific social network ontologies.
- *Lack of neutral user friendly tools.* Along FOAF project lifetime many projects have been developed, but almost none of them was oriented to support end users in editing and keeping updated their FOAF profiles. In this way, also the people that created a FOAF profile had to manually edit the RDF code describing its social network, discouraging users in keeping up-to-date their profiles or even worst creating invalid RDF documents.

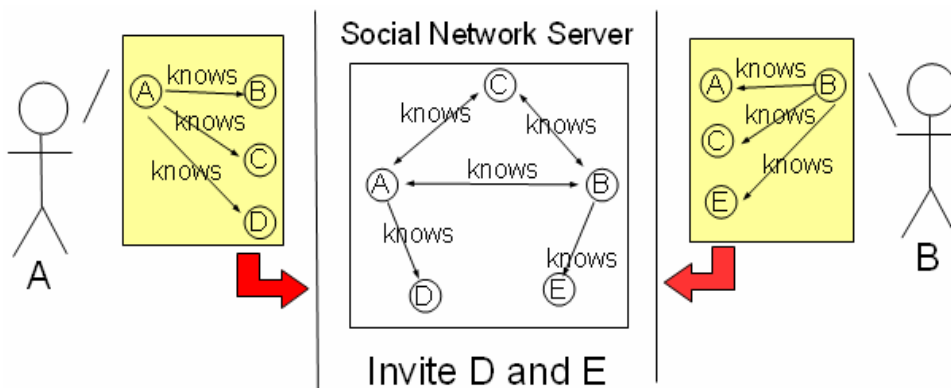
Similar issues and limitation of FOAF and other semantic web-based social network systems can be found in Sapkota et al. (2005).

We are convinced that the use of shared globally unique identifiers, combined with the use of a shared ontology, would solve the problem of automated social network integration. It is clear that a system like this must rely on a solid third-party identification system. Indeed, different users should be enabled to identify equally the same person across different private social network pieces of information. OKKAM ENS is the more general freely available naming service. Thereby, our solution for social network portability is the integration into FOAF personal profiles of globally unique identifiers for person provided through ENS interface. The owner of the FOAF social network profile should then be enabled to keep access control over the described information. In order to support users in managing FOAF profiles we defined a web-based rich internet application named foaf-O-matic. Such application is briefly described in Section 4.1.

It is important to notice how our approach is different from other identification approaches, such as the Single Sign-On. Our approach allows the global identification of persons in a completely application independent fashion. In principle, any application is enabled to recognise whether two persons have been virtually identified in the same way. This is not possible using the Single Sign-On approach, since it only allows a specific user to use their identifiers for accessing some social network service application.

Joining a social network application should consist of a simple operation: submit the social network profile to the social network application and accept its policy in terms of privacy management (e.g., information disclosure) and other related issues. The social network application should then recognise automatically all the people contained in the submitted social network profile that are already using the application and build the relative connections (see Figure 3). The submission could be realised by uploading the social network profile or simply by providing a link to the document. The second option has the advantage that the user could keep up-to-date its social profile only in one location and any modification would affect any social network application.

**Figure 3** Social networks smooth integration (see online version for colours)



As previously mentioned, an undesirable side effect of this approach is that everyone would be able to access the exposed social network information. A way to avoid this problem could be the creation of a repository with access control policies for the social network information. The social network information owner would then enable the social network application in accessing the information in a transaction on the basis of some secure access protocol (e.g., following the model of online payments).

It is important to notice that from this point on the only responsible for the security of the personal data 'submitted' is the social network application, where the social network activities are managed. From this perspective, nothing changes respect to the current situation. Indeed social network applications are still responsible to providing adequate and safe environment to host social network activities and handle personal information. Using a metaphor, the decentralised social network management is about opening 'safe communication channels', virtually enabling frictionless moving of social network related information between 'walled gardens' rather than abolish them. For this reason, we believe that data security is part of the big picture, but it is not the main issue we are aiming to solve.

### *5.1 The proposed solution*

In this section we outline our solution for satisfying the requirements for the realisation of a concrete decentralised social network management according to the paradigm described above.

#### *5.1.1 Decentralised social network management*

The first element necessary to build a system for decentralised social network management is the definition of a proper ontology. The FOAF ontology is the perfect starting point for the development of a decentralised social network management as it provides the first basic set of data type and object properties enabling the definition of a social network. Furthermore, being OWL ontology, it can be easily extended in order to define any specific kind of relations between persons. For this reason we chose FOAF as a generic starting point for the implementation of our solution. However, as already mentioned previously, FOAF alone does not satisfy the expressiveness requirements of current social network applications. Thereby, future evolution will consider more specific or specialised ontologies.

The second element necessary to realise decentralised social network management is the definition of neutral and user friendly toolkit which supports the definition, editing and storage of social network documents. In this article we present a web-based tool for the definition and editing of FOAF profiles, named foaf-O-matic. Foaf-O-matic is a rich internet application based on Icesoft Icefaces 1.7.1 J2EE AJAX framework that enables users to manage their social network profiles (see <http://www.foaf-o-matic.org>). A user can load and edit the personal social network profile, and integrate globally unique identifiers retrieved through guided interaction with an ENS.

As it is possible to see in Figure 4, the application allows users to describe the profile of the primary person in the left part ('me'), and describe and edit a list of friends in the right one ('my friends'). Each friend, as the primary person, can be described using an arbitrary number of FOAF attributes (point 3 in Figure 4).

Figure 4 View of foaf-O-matic interface (see online version for colours)

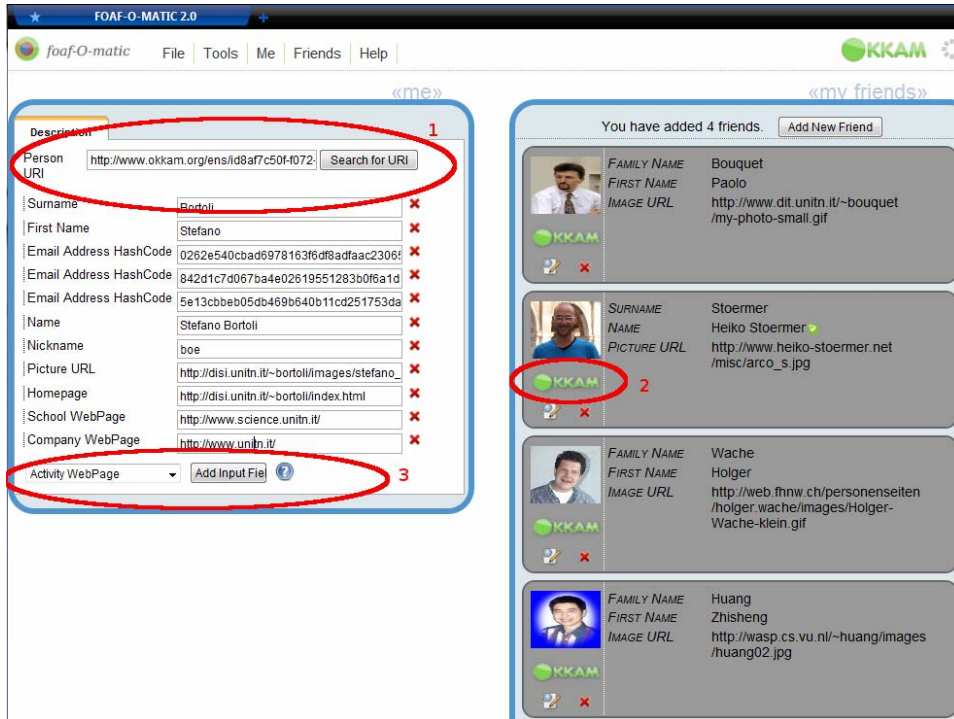
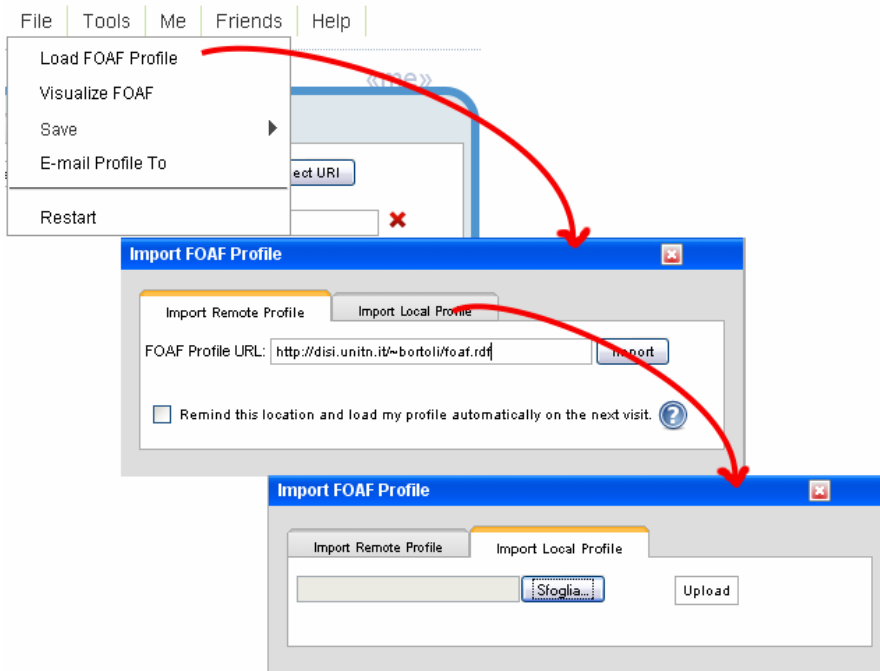


Figure 5 Load a FOAF profile (see online version for colours)



The FOAF ontology is constantly under evolution and in the FOAF community there is an ongoing discussion about the introduction of new attributes and properties, or the modification of existing ones. For this reason, we decided to make available for description only the attributes considered stable. This is done with the aim of generating RDF FOAF profiles that keep their compliance with the FOAF ontology as much as possible in time. The menu placed in the header of the application offers access to several functionalities. For example, it is possible to load a FOAF profile from local machine and remote location through the load FOAF profile panel (see Figure 5). If the loaded file contains a valid FOAF profile, all the information contained in it are loaded in the description forms as shown in Figure 4 and are ready for editing. The user can assign globally unique identifiers to the people described in the profile, add edit and remove friends' description.

Once the editing operations are complete, it is possible to save the FOAF profile on the local machine (see Figure 6) or post it on a remote location via SFTP connection (see Figure 7). The last version of the application supports also the option of sending via email a FOAF profile. In order to guarantee transparency, it is possible to visualise at any moment the RDF FOAF code corresponding to the current state of the social network profile (see Figure 8). In order to handle FOAF RDF loading and generation, foaf-O-matic relies on the widely used Jena libraries (<http://jena.sourceforge.net/>). Notice that the FOAF profile defined through foaf-O-matic does not include any attribute presenting a readable email address, but replaces it with an encrypted SHA1 checksum as the first foaf-a-matic (<http://www.ldodds.com/foaf/foaf-a-matic>) developed by Leigh Doods.

**Figure 6** Save FOAF profile on local machine (see online version for colours)

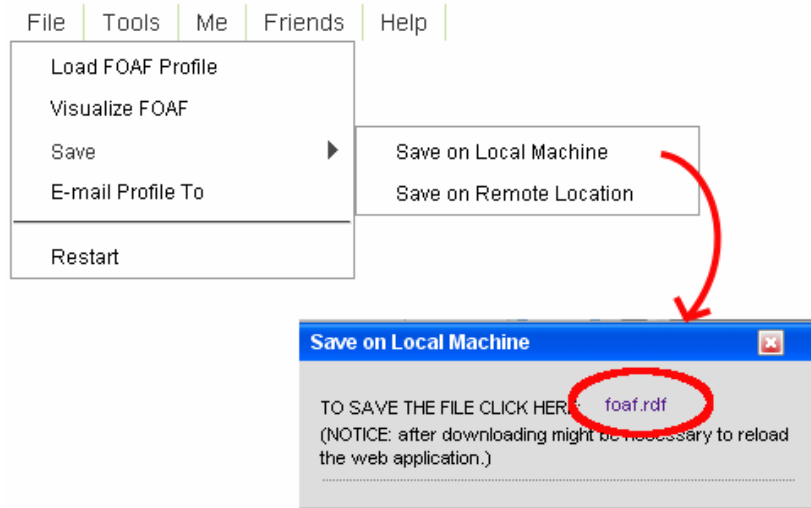


Figure 7 Post a FOAF profile on a remote location (see online version for colours)

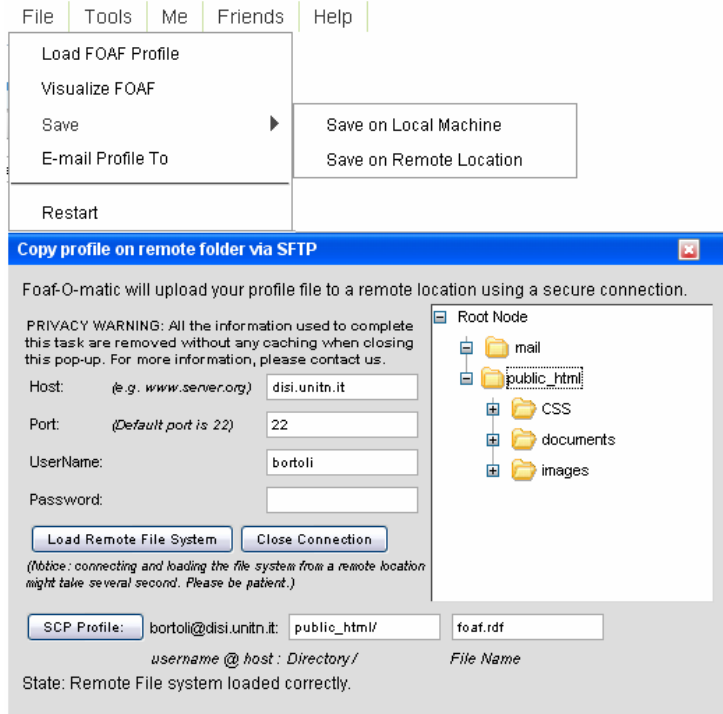
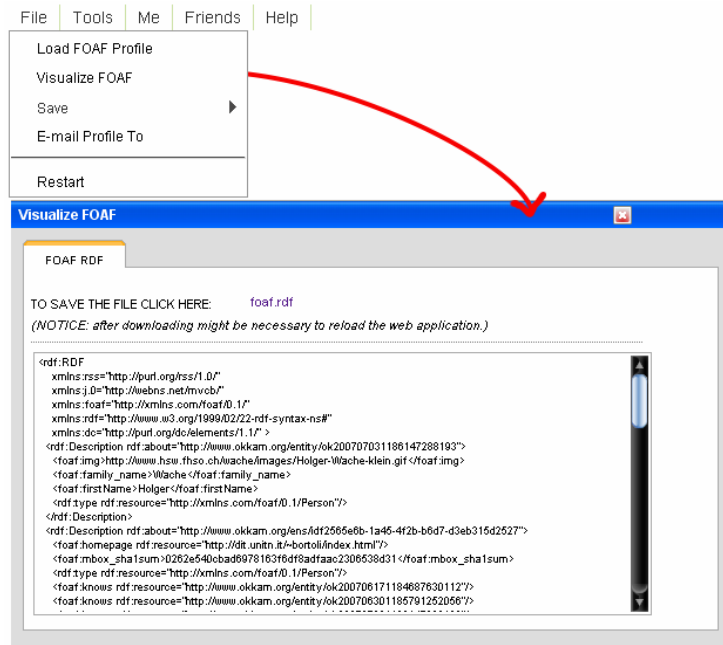
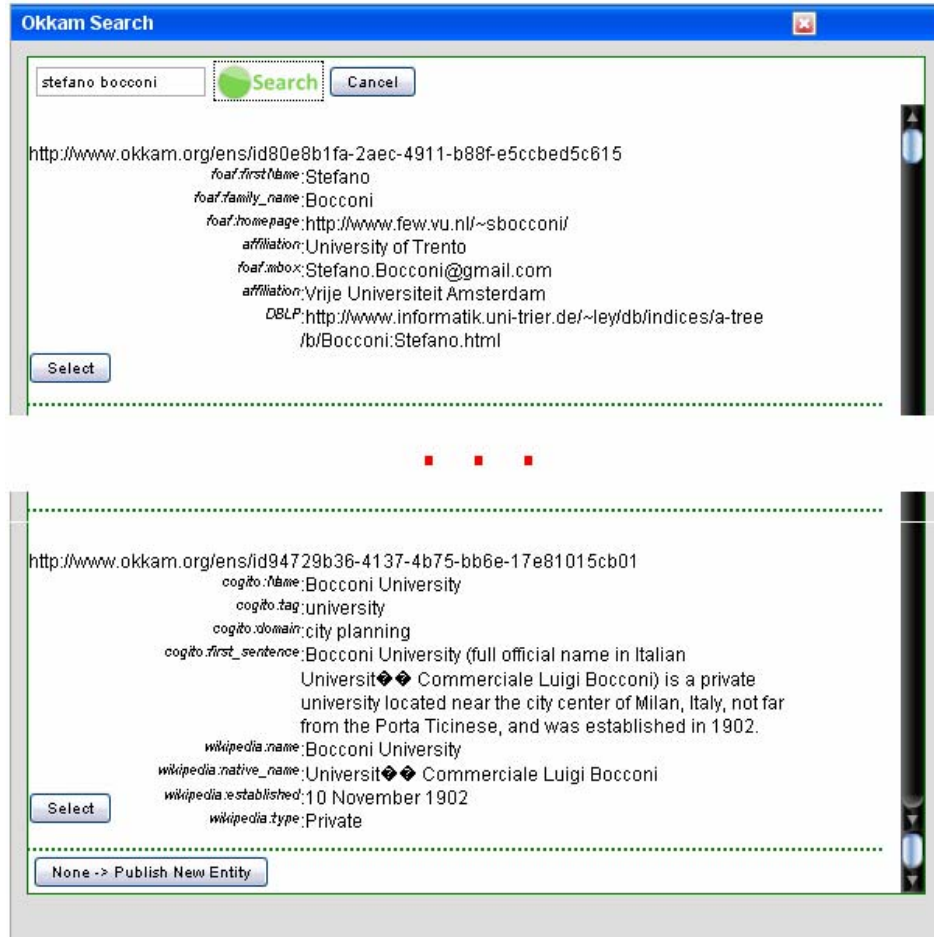


Figure 8 Visualise FOAF profile RDF (see online version for colours)





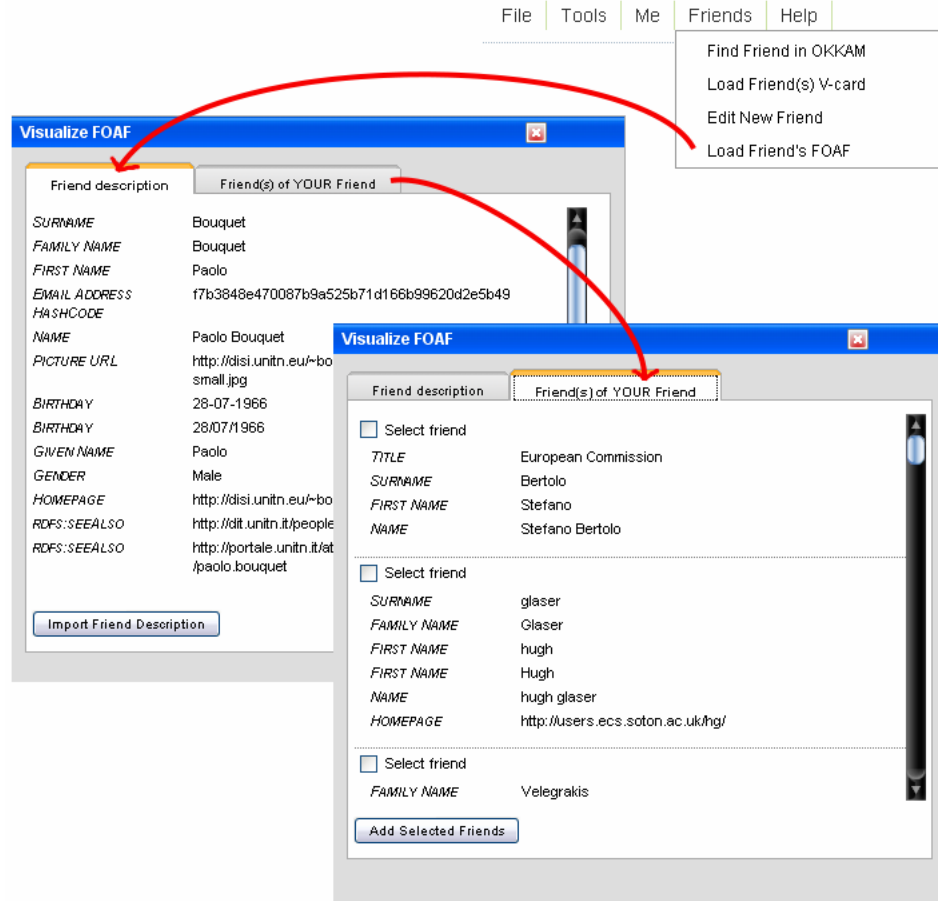
**Figure 9** Okkam search panel with query example (see online version for colours)

In order to speed up the social network definition process, it is possible to add descriptions of friends by loading their vCard (or list of vCards), or their FOAF profile. Importing person descriptions from vCards required a mapping between the vCard schema (<http://www.w3.org/TR/vcard-rdf>) and the FOAF ontology (<http://xmlns.com/foaf/spec/>). This was necessary because only few vCard attributes are suitable to be used in a FOAF profile. When user imports the FOAF profile of a friend, it is possible to the import also of 'friends of a friend' descriptions. Namely, a user can load the FOAF profile of a friend (if publically available) and add the description of the common friends to his/her own profile (see Figure 10).

Using the above tool, a social network user can load a social network profile, edit the personal information where necessary, and manage the definition of social relations. Every person mentioned in the profile can be identified with a globally unique identifier retrieved through the ENS. This operation is very simple: the user loads the ENS interaction panel, searches for the entity using some keyword describing the person (e.g., name and surname) and a list of candidate entity matching the query is displayed (see Figure 9). If the entity description for the searched person is among the candidate match

returned, the user can select the correct entity. This way, the identifier of the friend is integrated in the user FOAF profile of the user. If the searched entity description is not in the ENS, the user can then chose to start another procedure which will lead to the creation of a new entity description in the ENS. Other users wanting to refer to the same person will then be able to reuse the same identifier in their social network profiles.

**Figure 10** Import friends of a friend (see online version for colours)



As previously mentioned, the usage of a shared ontology is essential, but not sufficient to guarantee automated social network integration. In order to make this fundamental task possible, our solution relies on the ENS as source of shared, globally unique identifiers used to identify persons across social network profiles. For this reason, foaf-O-matic presents features enabling the user to interact with the ENS and embed globally unique identifiers in the social network FOAF profiles. The red ellipse marked with 1 in Figure 4 highlights the globally unique identifier retrieved through ENS and assigned to the primary person. When a friend is globally unique identified in the profile, the description is marked with a special logo as highlighted in point 2 in Figure 4.

Foaf-O-matic is in its 3<sup>rd</sup> version and is under constant evolution, in order to provide more services to support users in managing their social network in a decentralised and ‘neutral’ environment. Foaf-O-matic alone does not provide a complete solution to realise the decentralised social network paradigm described in Section 5, but it’s a first step towards the definition of an automatic decentralised social network integration system.

Following our approach, we plan to implement a web-based repository for social network profiles. This repository will be easy to integrate with foaf-O-matic and be the first complete toolkit enabling decentralised social network management.

### 5.1.2 Access control

In order to manage access control over the social network stored in the social network repository, we are considering endorsing the OAuth protocol (OAuth Core Workgroup, 2009). OAuth is an open protocol, initiated by Blaine Cook and Chris Messina, to allow secure API authorisation in a standard method for desktop, mobile and web applications. OAuth would allow a social network service provider to publish and interact with protected data stored on the social network repository, and would allow social network users to give access to their data to social network service provider while protecting their repository account credentials. In other words, OAuth allows a user to grant access to their information on one site (the social network profile repository), to another site (the social network service provider), without sharing all of their identity. OAuth is the result of the combination of other proprietary authentication protocols as:

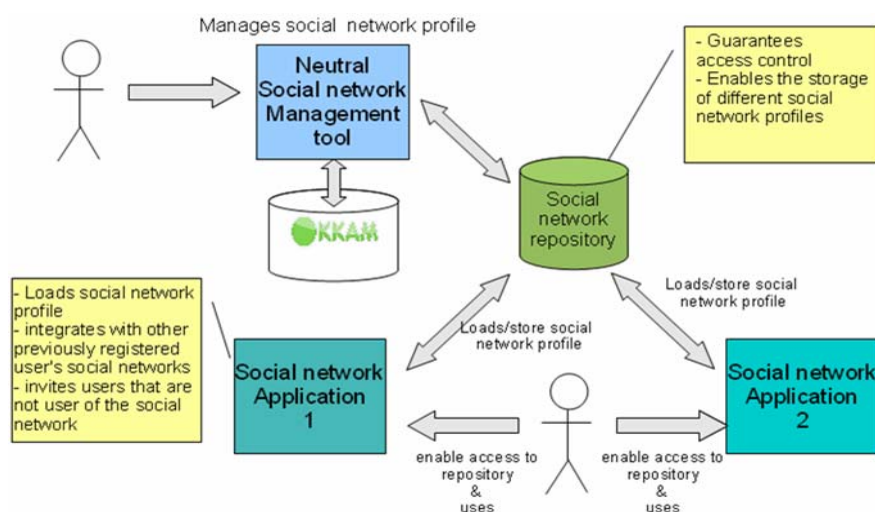
- Google AuthSub (<http://code.google.com/apis/accounts/AuthForWebApps.html>)
- AOL OpenAuth (<http://dev.aol.com/openauth>)
- Yahoo BBAuth (<http://developer.yahoo.com/auth/>)
- Upcoming API (<http://upcoming.yahoo.com/services/api/>)
- Flickr API (<http://www.flickr.com/services/api/>)
- Amazon Web Services API (<http://aws.amazon.com/>).

A graphical view over this possible solution is presented in Figure 11. At the top of the figure, we have the user interacting with a neutral social network management tool, editing a social network profile that integrates globally unique identifiers provided through an ENS. The neutral social network management tool then supports the interaction with a social network repository that can handle storage and access control of social network profiles. When the user wants to join a social network application, she must simply enable the latter to access her social network profile (bottom part of the picture). The social network application automatically integrates the loaded social network relying on the globally unique identifiers and, if the user requires it, invites all the friends that are not yet using the selected social network application.

Social network repositories are meant simply to provide storage and access control via network to personal data that, in principle, could also be maintained on the personal computer of the user. We don’t foresee these repositories as an active part of the social network profile management; we rather consider them some kind of remote storage unit providing adequate access control functionalities in pure web 2.0 spirit. For this reason, at the current state, the social network profile repositories are not meant to interact with

each other, or provide any other kind of intelligent service. Social network profile repositories should provide the same level of reliability, neutrality and verifiability as other existing applications (such as web-based remote storage application and web mail).

**Figure 11** A schematic view of a potential solution for decentralised social network management (see online version for colours)



## 6 Conclusions

In this paper, we presented an overview of the currently adopted approaches for enabling social networking. We also argued in favour of web-based communities that are decentralised and independent from social network web applications. We pointed out the limitations of the existing systems, and introduced a novel paradigm for the management of social networks in a decentralised and neutral environment.

The work presented in this paper is the first step towards decentralised social network management, since it offers a viable solution to a central problem: automatic social network integration based on globally unique identifiers. Foaf-O-matic is the first tool enabling the management of personal social network profiles outside any social networking web applications, since it relies on a machine readable standard vocabulary, and gives the possibility of uniformly identifying individuals across different social networks. User feedback showed that the current graphical interface of foaf-O-matic could be improved in terms of usability. Following this direction, we plan to work on the definition of a more user-friendly interface, which will guide users in managing their social network profiles.

We propose foaf-O-matic as a solution for two of the problems that affect the realisation of decentralised social networking management, namely, the lack of neutral user-friendly tools and the weakness of the identification system adopted by FOAF. In order to complete the realisation of the proposed paradigm, further steps are necessary, such as the definition of shared social network ontology. An extension of FOAF would be a promising direction to follow in order to accomplish this goal. This ontology should

present the attributes' metadata enabling the definition of access control policies, and therefore, limiting information disclosure. Another important step is the realisation of a social network profiles repository, with suitable access control policies. This way, users will be able to securely submit their social network profiles to the web applications of their choice.

## References

- Allsopp, J. (2007) *Microformats: Empowering Your Markup for Web 2.0*, Berkeley: Friends of ED, New York, USA.
- Andrews, P. (2003) *Social Networking Beginning to take Shape on the Web*, available at <http://community.seattletimes.nwsour.com/archive/?date=20031229&slug=paul29> (accessed on 14 September 2009).
- Bazzanella, B. et al. (2008) 'Towards a general entity representation model', *Proceedings of the 5th Workshop on Semantic Web Applications and Perspectives (SWAP2008)*, December, Rome, Italy.
- Bizannes, E. et al. (2009) *DataPortability: Vision and Mission*, available at <http://www.dataportability.org/> (accessed on 14 September 2009).
- Bojars, U. et al. (2008) 'Social network and data portability using semantic web technologies', *Proceedings of the 2nd Workshop on Social Aspects of the Web (SAW 2008)*, 6–7 May, Innsbruck, Austria.
- Bortoli, S. et al. (2007) 'Foaf-O-matic – solving the identity problem in the FOAF network', *Proceeding of SWAP 2007 – Fourth Italian Semantic Web Workshop*, 18–20 December, Bari Italy.
- Bouquet, P. et al. (2008) 'Entity name system: the backbone of an open and scalable web of data', *Proceedings of the IEEE International Conference on Semantic Computing, ICSC 2008*, August, pp.554–561, number CSS-ICSC 2008–4–28–25, IEEE Computer Society.
- Bouquet, P. et al. (2009) 'A conceptual model for a web-scale entity name system', *Proceedings of the 4th Annual Asian Semantic Web Conference (ASWC2009)*, December, Shanghai, China.
- Breslin, J.G. et al. (2005) 'Towards semantically-interlinked online communities', *Proceedings of the 2nd European Semantic Web Conference, ESWC 2005*, 29 May–1 June, Heraklion, Crete, Greece.
- Brickely, D. (2000) *About the FOAF Project*, available at <http://www.foaf-project.org/> (accessed on 14 September 2009).
- Brickely, D. and Miller, L. (2004) *FOAF Vocabulary Specification*, September, Namespace Document 2, FOAF Project, available at <http://xmlns.com/foaf/0.1/>.
- Euzenat, J. and Shvaiko, P. (2007) *Ontology Matching*, Springer-Verlag, Berlin Heidelberg, Germany.
- Freeman, L.C. (2006) *The Development of Social Network Analysis: A study in Sociology of Science*, Empirical Press, Vancouver USA.
- Google Press (2007) *Google Launches OpenSocial to Spread Social Applications across the Web*, available at <http://www.google.com/intl/en/press/pressrel/opensocial.html> (accessed on 14 September 2009).
- Kistner, J. (2007) *The 4Cs of the Decentralized Social Network*, available at <http://www.metafluence.com/the-4cs-of-the-decentralized-social-network/> (accessed on 1 January 2010).
- Mostarda, M. et al. (2008) 'Towards an OpenID-based solution to the social network interoperability problem', *W3C Workshop on the Future of Social Networking*, 15–16 January 2009, Barcelona, Spain.

- O'Reilly, T. (2007) *The Social Network Operating System*, available at [http://radar.oreilly.com/archives/2007/10/social\\_network\\_operating\\_system.html](http://radar.oreilly.com/archives/2007/10/social_network_operating_system.html) (accessed on 14 September 2009).
- OAuth Core Workgroup (2009) *OAuth Core 1.0 Revision A*, available at <http://oauth.net/core/1.0a> (accessed on 14 September 2009).
- Ramakrishnan, R. and Tomkins, A. (2007) *Toward a PeopleWeb*, Vol. 40, No. 8, pp.63–72, IEEE Computer Society Press.
- Sapkota, B. et al. (2005) 'SiFo-Peers: a social FOAF based peer-to-peer network', *Communications of the IIMA*, Vol. 5, No. 4, pp.81–90.
- Seok-Won, S. et al. (2009) *The Architecture and Implement of a Decentralized Social Networking Platform*, Stanford University, Stanford, CA 94305, available at <http://prpl.stanford.edu/>.
- Tungare, M. (2007) *The Case for Decentralized Social Networks*, available at <http://manas.tungare.name/blog/the-case-for-decentralized-social-networks/> (accessed on 1 January 2010).
- World Wide Web Consortium (1999) 'Resource description framework RDF model and syntax specification', *W3C Proposed Recommendation*, 05 January, available at <http://www.w3.org/TR/PR-rdf-syntax/> (accessed on 14 September 2009).
- Yeung, C-M.A. et al. (2008) 'Decentralization: the future of online social networking', *W3C Workshop on the Future of Social Networking*, 15–16 January 2009, Barcelona.