UNIVERSITY OF
BATH

Citation for published version:
Kelly, B & Dodds, L 2005, 'Using FOAF to Support Community Building', Upgrade Journal, vol. VI, no. 6.

Publication date:
2005

Link to publication

University of Bath

Download date: 12. May. 2019

Monograph of next issue (December 2005)
The Semantic Web
(The full schedule of UPGRADE is available at our website)

Using FOAF to Support Community Building

Brian Kelly and Leigh Dodds

The Semantic Web seeks to build a global distributed database through the integration of data from independent communities without the requirement for prior agreement on the structure of this data. This basic concept can be applied to the development of online communities. This paper outlines the potential for the FOAF (Friend Of A Friend) Semantic Web technology to ensure that resources can be defined in a way that promotes their ability for being shared with third-parties with a minimum of integration effort. The paper outlines FOAF's potential for community-building in conferences.

Keywords: Community Building, FOAF, Friend Of A Friend, Semantic Web.

1 Introduction

There is increasing potential for Web-based systems to support collaboration in many areas including education, business and social activities. As might be expected at this early phase of development we are seeing many products being developed. The wide range of different approaches used to construct collaborative tools is also evident, covering the well-established use of email, through to bulletin boards, instant messaging, video-conferencing, virtual conferencing, online voting, annotation, blogging, etc.

The dangers of application lock-in are well-known. The dangers apply in both commercial and open source software environments: in the commercial marketplace companies go out-of-business or are taken-over resulting in products becoming unavailable or unsupported; open source projects have similar longevity issues, with projects becoming moribund, unfashionable or superceded by other more successful solutions.

The data created within collaborative systems of all kinds has a potential to be of long term interest; indeed in the long term it is the data and not the individual products that are likely to be of real interest. We can envisage a need to ensure that not only message archives can be ported across collaborative systems but also richer forms of data, such as voting results, calendar entries, etc. Even in the short term there are advantages to be gained from easing data migration between systems, avoiding the need for end-users especially to reinvest effort in entering the same data into multiple systems. There is therefore a need to address the issues of migration of data across collaborative systems.

2 About The Semantic Web

The main advantage of XML (eXtensible Markup Language) is that developers can rapidly devise new document types that, because of their adherence to a basic syntax, can be easily manipulated in a number of standard ways. Beyond basic parsing these manipulations are described by the larger "XML family" of specifications that define standardized methods of transformation, inclusion, linking etc. With the addition of standard programmatic interfaces for manipulating XML data, this results in a great improvement in interoperability, allowing much easier processing and movement of data between systems. The end result has been that XML has been tremendously successful and has penetrated into many different kinds of application.

However at a slightly higher level XML adds little to the ability of applications to interpret or combine the data being transferred: XML applications have to be extended each time a new document type must be processed. Building knowledge of a wide range of XML formats into an application can be very time-consuming; only in the most trivial of cases can simple transformations be applied. This results in a continuing investment in application integration. One solution to this problem is to define a standard document type for interchange. However consensus-building is always a time-consuming process and such a solution only works for well-defined communities. For a looser community discovering the common ground ripe for standardization is much harder, typically resulting in extensions or variants in a format which mitigates some of the sought-after benefits.

The Semantic Web effort is aimed at tackling these integration efforts head-on, by building an infrastructure which greatly facilitates the merging of disparate data sets, allowing the creation of what might be regarded as a global "web of data". The key to this infrastructure is the Resource Description Framework (RDF) [6]. While RDF is an XML application, it emphasizes the data model over syntax: while

Brian Kelly provides advice and support to the UK Higher and Further Education Communities and the museums, libraries and archives sector in the area of the Web. He is based in UKOLN - a national centre of excellence in digital information management, based at the University of Bath (UK). <B.Kelly@ukoln.ac.uk>

Leigh Dodds is an Engineering Manager employed by Ingenta where he is responsible for the ongoing development of the IngentaConnect Web site, a large aggregation of academic research content. He has been developing with Java, XML, and Semantic Web technologies for a number of years and has released several small open source applications and APIs. He has also contributed code and documentation to several open source projects. <leigh@ldodds.com>

XML may be said to provide benefits by defining a common syntax, RDF derives its benefits from mandating a common underling model for representing data.

The RDF model is graph-based, containing resources (nodes) that are identified by a URI (typically a URL) and statements about them (arcs). The reliance on URIs means that RDF is intimately tied into the web infrastructure, while its mathematical underpinnings give it a rigorous definition. Although its basic elements are quite straight-forward, RDF can be a difficult concept to grasp and interested readers are directed to the RDF FAQ [7] for more information.

3 FOAF

Due to perceptions about complexity many Semantic Web applications are still being developed within the Web research community [8] with a slow but growing penetration into other areas; growth has certainly been nothing like as rapid as that of XML. One Semantic Web application noteworthy because of its active grass-roots development and user-centric focus is FOAF (Friend Of A Friend).

The core of FOAF is simple: provide a RDF vocabulary for capturing metadata about people. One interesting piece of metadata is a person's relationships to others; it is from this aspect that FOAF derives its name: "Friend-Of-a-Friend".

FOAF then, captures data pertinent to social networking applications and studies. Social networks are a current hot topic, with many applications appearing that are intended to allow people to mine their social networks for business and employment contacts, dating and friendship opportunities, product and reading recommendations, etc.

As well as directly relating people, FOAF also provides a vocabulary for describing a persons interests, their membership within groups of people and their project involvement. Other terms in the vocabulary allow statements such as "I made this" (e.g. declaring the author of a document) or "I own this" (e.g. items in a book or music collection). As a whole this provides a very rich set of metadata elements, enabling many different kinds of application. As an RDF vocabulary FOAF can be easily integrated with other RDF vocabularies allowing for example, the publishing or music community to define the optimum way of describing a book, album or other creative work, while still allowing this data to be related to people (the owner or creator).

One key premise of FOAF, and of RDF generally, is that data is inherently distributed across the Internet rather than centralized in a single database. There is no single source of FOAF data: individuals maintain their own FOAF descriptions, either directly or by using tools provided by their community. This data is then harvested using indexing applications that share many similarities to the web crawlers employed by Internet search engines. FOAF provides a facility for linking one FOAF document to another in the same way that HTML documents are linked across the Web; this has resulted in a web of interlinked FOAF data sources.

3.1 FOAF Applications

To date FOAF application development has centered on FOAF authoring tools, e.g. the FOAF-a-Matic [1] as well as tools capable of viewing and browsing FOAF data, e.g. FOAF Explorer [3] which generates Web pages from FOAF data allowing a user to read a person's self-description and browse to related information such as project homepages, lists of interests, or information about their friends and collaborators. One of the most interesting FOAF applications is FOAFNaut [4] which provides a visualization of the social network described by the distributed FOAF data set. This visualization is displayed in the browser using SVG (Scalable Vector Graphics) creating an easily navigable graphical view of the relationships between people. This social network browser is made possible from the metadata that people have included in their FOAF descriptions.

4 Using FOAF to Support Community Building

While FOAF has already generated interest from a number of existing online communities, including the Blogging community, the authors are interested in the ability for FOAF to foster community building. Specifically, the authors are exploring the use of FOAF to support community building amongst conference delegates by allowing delegates to record their past and intended attendance at conferences, along with their interests and other pertinent data. As well as providing an application to help delegates create the required metadata, the authors will produce supporting tools for exploration of this data allowing, for example, delegates with similar interests to find one another.

One of the authors (Kelly) organizes the Institutional Web Management Workshop, which has been held annually since 1997. The workshop provides a forum for members of Web management teams within UK Universities, and attracts over 150 delegates, many of whom are regular participants. For the forthcoming workshop it is intended to provide a tool for the delegates that will allow them to record their attendance at Institutional Web Management Workshop events, as well as other conferences; their interests; contact
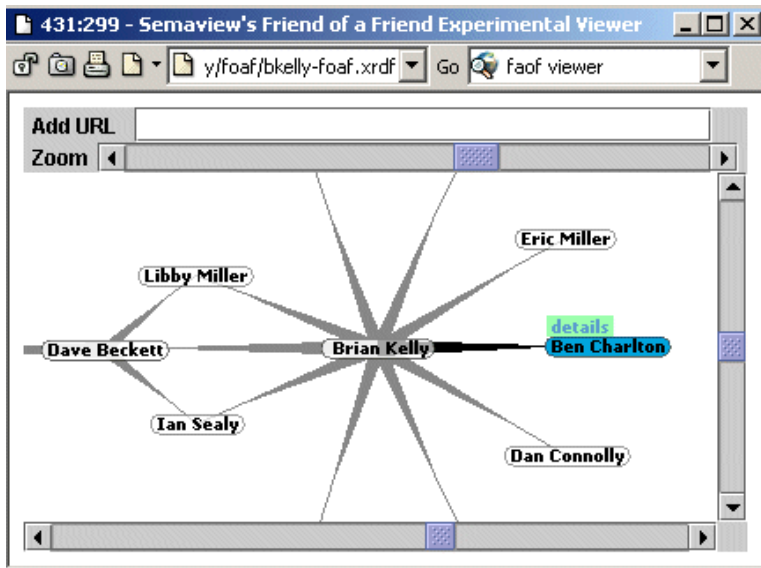

Figure 1: FOAF Authoring Tool.

Figure 2: Example of A AFOAF viewer.

details; and other appropriate information, e.g. a photograph, job status, institutional affiliation, role at the event, professional interests and links to people they know.

Delegates will be able to use the tool prior to the event and workshop speakers and organizers in particular will be encouraged to use the tool to ensure a useful data set is available. The authoring tool is being adapted from the existing application [1] tailoring the interface to support this specific use as shown in Figure 1.

Rather than manage the data centrally, delegates will publish their own FOAF descriptions which can then be harvested to build a number of potentially useful services.

In particular the authors will demonstrate how existing FOAF visualization tools such as FOAFnaut and FOAF Explorer can be used to explore this Semantic Web space. Additionally a custom browsing interface will be created allowing users to answer queries such as (a) show all delegates at the event; (b) show delegates from a particular institution and; (c) show delegates with an interest in information architecture.

An example of a FOAF viewer is shown in Figure 2. It should be noted that this diagram illustrates very basic use of FOAF – a richer application would allow, for example, the user to access information by events attended and allow events attended by an individual to be displayed.

Once delegates have seen the potential we hope that this will provide motivation to provide information on attendance at other events. Once a proof-of-concept has been developed we can then justify the development of additional tools to complement the browsing interface.

5 What Does This Give Us?

It could be argued that this application outlined above could be implemented using conventional collaborative tools. However this fails to acknowledge the main point of an RDF-based approach – ensuring that the underlying data can be integrated with current and future applications. For exam-

ple delegates with existing FOAF descriptions can easily supplement their existing metadata with that collected by this application. This can be achieved without any agreement with, or notification to, other consumers of FOAF data. Similarly other application developers interested in using the additional conference related metadata generated by this study can easily develop or extend applications to process it.

6 The Next Steps

We hope that the work described in the paper will provide a valuable tool. However it is recognised that the work is in early stages. Further work is needed in developing FOAF viewers which will exploit the conference data described in this paper; developing a range of addition FOAF creation tools, which can build on the Java tool described in this paper and engaging with user communities in encouraging use of the tools and ideas described. We hope that this approach described in this paper will help the delegates to gain a glimpse of the potential of Semantic Web applications such as FOAF and encourage further experimentation. We welcome feedback from the community and invite those who would like to be involved in further development work to contact the authors.

Acknowledgements

References
[1] [FOAF-a-Matic, L. Dodds. <http://www.ldodds.com/foaf/foaf-a-matic>.
[2] FOAF, <http://www.foaf-project.org>.
[3] FOAF Explorer, M. Frederiksen. <http://xml.mfd-consult.dk/foaf/explorer/>.
[4] FOAFNaut, J. Ley. <http://www.foafnaut.org>
[5] Web Based Communities 2004 Conference, IADIS. <http://www.iadis.org/wbc2004/>
[6] Resource Description Framework (RDF), W3C. <http://www.w3c.org/RDF/>
[7] RDF FAQ, W3C. <http://www.w3.org/RDF/FAQ>.
[8] RDF Projects and Applications, W3C. <http://www.w3.org/RDF/#projects>.