

Social Networking on the Semantic Web

Tim Finin, Li Ding and Lina Zou
University of Maryland, Baltimore County
Baltimore MD USA

The Semantic Web promised to enable a new generation of intelligent applications by providing programs and software agents with rich and effective ways to share information and knowledge. The Semantic Web allows people and groups to define sharable ontologies – collections of classes, properties and objects – with well defined and unambiguous meaning. These ontologies permits computer programs to read, publish and exchange information and knowledge, enhancing interoperability, cooperation, and service-oriented architectures. While we are early in the development and exploration of the Semantic Web, there is evidence that the concepts and technologies are being adopted and used. One of the most widely used current ontologies is FOAF, or "Friend of a Friend". The FOAF ontology specifies a vocabulary that can be used to define, exchange and search for social information – information describing people, their attributes and their relationships with others. We will briefly review the concepts underlying social networks and the Semantic Web and then describe how they are embodied in FOAF and being used today on the Semantic Web.

Draft. To appear, Learning Organization Journal, special issue on Ubiquitous Business Intelligence, Miltiadis Lytras et al, 2005. Author contact: Professor Tim Finin, Computer Science and Electrical Engineering University of Maryland, Baltimore County, 1000 Hilltop Circle, Baltimore MD 21250. finin@umbc.edu, voice: +1.410.455.3522, fax: +1.410.455.3969

1 Introduction

Social networks are explicit representations of the relationships between individuals and groups in a community. In the abstract, these networks are just simple graphs with nodes for the people and groups and links for the relationships. In practice, the links can encode all kinds of relationships – familial, friendship, professional or organizational. Social network theory, the study of such social networks, has developed techniques found useful in many fields, including sociology, anthropology, psychology and organizational studies. Social network analysis has been used, for example, to represent and analyze the organization structure of employees in a business unit, identify key individuals, and suggest structural changes to improve unit performance.

Virtual or online communities are groups of people connected through the Internet and other information technologies. These have become an important part of modern society and contribute to life in many contexts - social, educational, political and business. The communication technologies and infrastructures used to support virtual communities have evolved with the Internet and include electronic mailing lists, bulletin boards, usenet, IRC, Wikis, and blogs. Virtual communities built on social network structures began appearing in 2002 and have become among most popular Web-based applications. Such sites allow individuals to publish personal information in a semi-structured form and to define links to other members with whom they have relationships of various kinds. Current examples include Friendster, LinkedIn, Tribe.net, and Orkut. Other web-based virtual communities have successfully combined social networking with various in-

terests, such as photography (Flickr.com), film (Netflix.com), personal blogging (Myspace.com) and dating (Thefacebook.com).

Several of these social network based virtual communities have begun to publish members' public profile information, including social links, using the Semantic Web language RDF. Most use the RDF vocabulary defined by the "Friend of a Friend" (FOAF) ontology augmented with new terms as needed. The use of a widely known, non-proprietary, shared ontology for this information enables interoperability among these systems. More importantly, RDF was designed as a data sharing standard privileging extensibility. Individual systems can extend the vocabulary as needed without interfering with the ability to combine and integrate information. This approach opens up many possibilities for information integration, aggregation and fusion on the Web.

1.1 The Semantic Web and ontologies

The Semantic Web idea emerged from the confluence of several communities - artificial intelligence, hypertext, Web developers - and so there are a number of ways to appreciate its motivation and goals. Perhaps the easiest for one who does not belong to any of those communities is to consider that much of what we want to know (that is actually known) is available on the web. Thus the Web is, potentially, a great resource for software agents, which can be programmed to extract and fuse information from multiple, heterogeneous sources in response to a query.

However, extracting meaning from text is a very challenging task for computer programs. While progress is being made, a robust solution is decades, if not generations away. So the Semantic Web is an approach to encoding and publishing information in ways that makes it easier for computers to understand, thus making the Web agent-friendly. What do we mean by "making it easier for computers to understand?" On the Semantic Web, we mean: through recourse to ontologies, formal descriptions of particular domains.

Ontology is the branch of philosophy that seeks to answer the question "what is there? ." In computer science, an ontology is a formal conceptualization of a domain. Typically, it specifies the classes of objects that exist, the relationships amongst those classes, the possible relationships amongst instances of the classes, and constraints over those instances. An ontology also defines terms denoting these classes and relationships as well as individual objects. Current web ontology languages, designed to encode information on and for the web, use the eXtensible Markup Language or XML both for specifying ontologies, and also for making assertions about the world using terms defined in ontologies. A Semantic Web page begins by listing (as URLs) the locations of the ontologies to be used, then goes on to use those ontologies to make assertions about datasets, human beings, items for sale, etc. An agent, on coming to such a page, can import the specified ontologies and use that information to understand the semantics of the ensuing assertions.

The World Wide Web Consortium (W3C) has developed standards to enable ontologies to be published on the Web as well as data and other assertions to be encoded using terms drawn from any published ontologies. These standards make it possible for programs and software agents to understand information published on the Web without the ambiguity and complex processing inherent in traditional unstructured forms (e.g., natural language) or rigidity and lack of flexibility inherent in structured representations (e.g., relational databases.)

The Resource Description Framework or RDF [**Error! Reference source not found.**] is a simple XML-based language to define computer-understandable vocabularies that people and

programs can use to describe things of interest, such as Web sites, newspaper articles, email messages, people, books, events, or Web services. RDF mimics human languages in that it allows one to introduce new terms (individuals, classes and properties) that are defined (partially, at least) in terms of existing terms. RDF Schema **[Error! Reference source not found.]** extends RDF by providing vocabulary to build logical object-oriented schema, including a simple typing system, sub-classes, sup-properties, inheritance, etc. The Web Ontology Language OWL **[Error! Reference source not found.]** supports advanced capabilities, such as logical inference and translating descriptions using different ontologies (e.g., mapping a location specified as a ZIP code to one using latitude and longitude.)

A problem in the effort to formalize (or “ontologize”) a domain is that there are typically many different ways of doing so. This is true whether the domain is in a science, or business-related or has to do with people and their relationships. Within a single discipline, there can be disagreement about how to describe the world. As well, disciplines overlap, and often look at the overlapping area from different points of view. One approach to the ontology heterogeneity problem is to create a global schema to serve as an Interlingua for human and software agents. One of the principles of the Semantic Web is that it should be based on the same open, decentralized and distributed approach that has made the World Wide Web successful. Anyone should be able to create, publish and use their own ontologies. Mechanisms are available to allow one to define mappings or translations of terms among ontologies. In the open and dynamic environment of the web, it is expected that the natural influences and forces of the market and “networking effect” will encourage coalescing to a smaller number of interoperable ontologies for a given domain. So the construction of a few global schemata is not the goal. Rather, we envision and are encouraging the development of a number of relatively small ontologies, some of which may overlap, and some of which may be in conflict.

1.2 The Friend of a Friend Ontology

The FOAF vocabulary includes classes and properties found useful to describe people online. Consider the following example, drawn from the FOAF Vocabulary Specification **[Error! Reference source not found.]** and encoded using the XML serialization for RDF.

```
<foaf:Person>
  <foaf:name>Dan Brickley</foaf:name>
  <foaf:mbox_sha1sum>241021fb0e6289f92815fc210f9e9137262c252e</foaf:mbox_sha1sum>
  <foaf:homepage rdf:resource="http://rdfweb.org/people/danbri/" />
  <foaf:img rdf:resource="http://rdfweb.org/people/danbri/mugshot/danbri-small.jpeg" />
</foaf:Person>
```

This example encodes the information that “there is a *foaf:Person* with a *foaf:name* property of ‘Dan Brickley’ and a *foaf:mbox_sha1sum* property of 24...52e; this person stands in a *foaf:homepage* relationship to a thing called *http://rdfweb.org/people/danbri/* and a *foaf:img* relationship to a thing called *http://rdfweb.org/people/danbri/mugshot/danbri-small.jpeg*.” FOAF defines 12 classes and 51 properties. The *foaf:knows* property is used to construct basic social networks, linking to instances of *foaf:Person*.

The FOAF vocabulary is simple, which has encouraged its adoption and use, and extensible, making it suitable to a wide range of uses. As our studies have shown, more than 150 different

properties have been defined for the *foaf:Person* class and nearly 500 have actually been used with instances of *foaf:Person*. One way to view this situation is that it represents undisciplined chaos and that the lack of any centralized authority or standard for terms suggests that nothing useful will come out of it. An alternate view is that communities will be able to select and use terms that are useful and those which are widely used be integrated into consensus ontologies. In this view the eventual result will be a relatively small number of widely used ontologies with mappings, as appropriate, between them. Less widely used terms, whether they are deprecated, or newly introduced, will remain on the edges.

Our investigation the most commonly used ontologies (see Table 1) confirms that, besides the meta-level ontologies (i.e. RDF, RDFS, DAML and OWL), one of the best populated ontology is FOAF (Friend-of-a-Friend) [**Error! Reference source not found.**]. In addition, representing personal information is also a popular theme in ontology engineering with more than a thousand RDF documents defining RDF terms containing the string ‘person’¹. The other well populated ontologies in Table 1 include DC (Dublin Core Element Set) [**Error! Reference source not found.**], which defines document metadata properties without domain/range qualification, and RSS (RDF Site Summary), which is “a lightweight multipurpose extensible metadata description and syndication format” for annotating websites [**Error! Reference source not found.**]. FOAF provides an RDF/XML vocabulary to describe personal information [**Error! Reference source not found.**], including name, mailbox, homepage URL, friends, and so on. FOAF documents then induces the “web of acquaintances” [**Error! Reference source not found.**] and thus an implicit trust network to support such applications as knowledge outsourcing [**Error! Reference source not found.**] and online communities [**Error! Reference source not found.**].

Table 1: Best populated ontologies (generated in April, 2005)

Prefix	Namespace URI	Documents populated
RDF	http://www.w3.org/1999/02/22-rdf-syntax-ns#	321,108
DC	http://purl.org/dc/elements/1.1/	238,346
RSS	http://purl.org/rss/1.0/	195,018
MCVB	http://webns.net/mvcb/	110,434
FOAF	http://xmlns.com/foaf/0.1/	79,226
RDFS	http://www.w3.org/2000/01/rdf-schema#	65,486

The advances in FOAF vocabulary and applications highlight several challenging issues. For example, how can one assemble a collection of FOAF documents to support Semantic Web research? What are the common patterns of connections among FOAF documents? What terms in FOAF vocabulary are the most frequently used? What is the potential of FOAF in enabling and enhancing the intelligence of Web-based information systems? The current FOAF literature

¹This is reported by our Swoogle (<http://swoogle.umbc.edu>), a RDF crawling and indexing engine [**Error! Reference source not found.**].

[Error! Reference source not found.; Error! Reference source not found., Error! Reference source not found.b, Error! Reference source not found.; Grimnes, Edwards, & Preece, 2004; **Error! Reference source not found.**] provides a vision and various models of how FOAF documents might be used to support Web-based information system under the assumption that FOAF documents are widely available. There is still a lack of an empirical investigation on the characteristics and structure of the growing body of millions of FOAF documents. This paper presents empirical results to answer the above questions based on a large collection (over 1.5 million) of real world FOAF documents harvested from the Web.

Our research on online FOAF profile documents consists of four steps: identification of FOAF documents, discovery of FOAF documents using software agents, extraction of person information, and fusion of person information based on the semantics of FOAF vocabulary. Using the statistics over this corpus, we describe the common properties and namespaces shared by the FOAF community. We hope that this analysis might help FOAF developers design and build better tools as well as inform novice FOAF users on how to create effective FOAF documents. Analyses of the social networks encoded in FOAF documents provide insight into some interesting structural patterns of the Semantic Web from the person perspective. The richness of profiles in FOAF documents allows us to further characterize social ties and identify friendship types.

Friendship networks connected by FOAF relationships can provide insights into features and patterns of social networks in the Semantic Web and advance the theories and models of social structures. Friendship networks in the physical world have been long studied in the social science. A well known example is Milgram's small-world phenomenon **[Error! Reference source not found.]** - the observation that everyone in the world can be reached through a short chain of social acquaintances. The concept gives rise to the famous phrase six degrees of separation, which has recently been applied to social network analysis in both physical and virtual environments, e.g. **[Error! Reference source not found.; Error! Reference source not found.]**. Social relationships have been derived from the contextual information or domain knowledge, e.g. co-citation relationship **[Error! Reference source not found.]**, indirectly using data mining techniques. In addition to social networks, the collection of FOAF documents can serve as valuable resource for Semantic Web research in the development and testing of trust models as well as trust propagation models **[Error! Reference source not found.]**.

As the first study along this line, this paper reflects the state of FOAF usage and identifies any potential problems to guide the future practice. It further contributes to the stabilization of individual terms in FOAF vocabulary. Using people as the bridge, FOAF can potentially link most of other kinds of things we describe in the Web, including documents they co-authored, research interest they shared, photos they shot together, and so on. Based on relationships represented in FOAF, we can identify online communities in a research area and even discover existing communities and the emergence of new communities. As the Semantic Web evolves, there will be opportunities to study social dynamics and apply the findings in this study to support Semantic Web applications.

The remainder of this article is organized as follows. Section two presents a review of the literature concerning FOAF vocabulary and social network analysis. Section three introduces a novel approach to building FOAF documents collection and analyzing the structure of friendship networks in the Semantic Web. Section four uses descriptive statistics and social network analysis to present findings on components of FOAF documents and structural relationships among person

profiles. Section five concludes with a discussion the findings of this study and their implications to the Semantic Web research and practice.

2 Background

2.1 FOAF Document

The most important component of a FOAF document is the *FOAF vocabulary*, which is identified by the namespace URI <http://xmlns.com/foaf/0.1/>. The FOAF vocabulary defines both classes (e.g., *foaf:Agent*, *foaf:Person*, and *foaf:Document*) and properties (e.g., *foaf:name*, *foaf:knows*, *foaf:interests*, and *foaf:mbox*) grounded in RDF semantics. In contrast to a fixed standard, the FOAF vocabulary is managed in an *open source* manner, i.e., it is not stable and is open for extension². Therefore, inconsistent FOAF vocabulary usage is expected across different FOAF documents.

The practical significance of FOAF to information creators and consumers can be illustrated with a variety of applications [**Error! Reference source not found.**, **Error! Reference source not found.**], which are summarized as follows. To information publishers, FOAF is useful by

- Managing communities by offering a basic expression for community membership. Many communities have proliferated on the Web, ranging from companies through professional organizations to social groups.
- Expressing identity by allowing unique user IDs across applications and services without compromising privacy. For example, the *foaf:mbox_sha1sum* property is the ASCII-encoded SHA1 hash of a mailbox URI (e.g., <mailto:finin@umbc.edu>). To ensure privacy, the encoding is a one-way mapping and cannot be trivially reverse-engineered.
- Indicating authorship. FOAF tools use digital signatures to associate an email address with a document. Specifically, OpenPGP is used, along with the namespace <http://xmlns.com/wot/0.1/> to denote concepts forming a “web of trust”. This associates a signature with the document itself and specifies a signature for the linked document as part of an *rdfs:seeAlso* link. Thus, authorship information can be expressed both inside and outside of the concerned documents.

FOAF supports information consumers by:

- Allowing provenance tracking and accountability [**Error! Reference source not found.**]. On the Web, the source of information is just as important as the information itself in judging its credibility. Provenance tracking RDF tools can tell where and when a piece of information is obtained. A practice common to the FOAF community is to attach the source URI to each RDF statement.
- Providing assistance to new entrants in a community. For example, people unfamiliar with a community can learn the structure and authority of a research area from the community’s FOAF files.

²The latest FOAF specification only lists one stable term – ‘homepage’ and leaves many others in ‘testing’ or ‘unstable’ stages.

- Locating people with common interests. Users tend to have interests and values similar to those they desire in others [**Error! Reference source not found.**]. Peer-to-peer relationships are an essential ingredient to collaboration, which is the driving force of online communities.
- Augmenting email filtering by prioritizing mail from trustable colleagues. Using the degree of trust derived from FOAF files, people can prioritize incoming email and thus filter out those with low trust values.

2.2 Social Networks on the Web

A social network consists of people or groups connected by a set of social relationships, such as friendship, co-working or information exchange [**Error! Reference source not found.**]. Determining structural properties of virtual communities is the most straightforward application of social network analysis (SNA). The underlying physical social network can be reflected in an online community. For example, Club Nexus [**Error! Reference source not found.**] is an online community serving over 2000 Stanford undergraduate and graduate students. Students can use Club Nexus to send email and invitations to events, post events, buy and sell goods, search and connect to people with similar interests, etc. Statistical analyses revealed that personalities and preferences of users mostly align with each other.

In addition to member relationship in online communities, SNA has been applied to many other types of social networks. For example, [**Error! Reference source not found.**] created, analyzed and visualized a network of known criminals and their relationships. Their analysis identifies various groups and subgroups, key individuals, and links between groups. Centrality can be detected using graph properties including degree (the number of direct links), betweenness (geodesics passing through), and closeness (sum of geodesics). Each of these indices is evidence for different individual roles: a high degree suggests leadership and high betweenness indicates a “gatekeeper”. This increased understanding enables law enforcement officers to target specific criminals, to disrupt criminal organizations, and to achieve higher rates of conviction.

Chen [**Error! Reference source not found.**] describes the development and application of visualization techniques allowing users to access and explore information in a digital library effectively and intuitively based on co-citation relationships. Salient semantic structures and citation patterns are extracted from several document collections using latent semantic indexing and pathfinder network scaling. Author co-citation patterns are visualized through a number of author co-citation maps highlighting important research areas in the field. This approach provides a means of transcending the boundaries of collections of documents and visualizing more profound patterns in terms of semantic structures and co-citation networks.

Link structure analyses and graph-theory have been applied to crawling the Web for virtual communities [**Error! Reference source not found.**]. The FOAF project takes the social networking aspect of the Web further [**Error! Reference source not found.**]. An interesting and powerful use of information in FOAF documents arises when data are aggregated and then cross-linked.

3 Discovering FOAF information on the Web

By running the SwoogleBot [Error! Reference source not found., Ding et al. 2005] Semantic Web crawler in conjunction with an agent that understands FOAF vocabulary, we collected 49,750 RDF documents containing 207,413 instances of *foaf:Person* during the first three months of 2005. We intentionally limited the dataset by collecting at most 50,000 documents from any single website and no documents from several large blog sites (e.g., *livejournal.com*).

3.1 Provenance of the data

Table 2 lists the five community websites with the most number of FOAF documents. We identify several different contexts in which this information is used: to describe blog authors, to describe virtual community members, or to annotate photographs.

Table 2: Community Websites

Host	Context	FOAF dataset	Swoogle discovered	Google site estimation
www.livejournal.com	blog	avoid	46,661	5,370,000
www.tribe.net	community	avoid	23,518	2,920,000
blog.livedoor.jp	blog	4,266	10,120	119,000
www.greatestjournal.com	blog	avoid	10,097	282,000
www.wasab.dk	annotation	4,910	8,434	73,700

Although community websites have contributed large numbers of FOAF instances, their regular structure also overwhelms the variety of vocabulary and structure introduced by people who construct and self-publish FOAF profiles. We adopted a simple heuristic applied to URLs, to recognize those from community websites. If there are a large number of URLs from a given site that differ only in a single URL argument, we classify them as automatically generated. Table 3 shows some extracted URL pattern³ for community websites.

Table 3: URL Patterns for Community Websites

URL pattern	example match (amount of matches)
same host and path, different query	http://www.boards.ie/network/foaf.php=? [QUERY] (2490)
same host and path, different query	http://www.boards.ie/network/foaf.php=? [QUERY] (2490)
same host, no query, path differs in one segment	http://journal.bad.lv/users/[USERNAME]/data/foaf (2548) http://blog.livedoor.jp/[USERNAME]/foaf.rdf (4242) http://swordfish.rdfweb.org/photos/genfiles/ilrt/[FILENAME](266)
same path, no query, host differs in first segment	http://[USERNAME].cocolog-nifty.com/foaf.rdf (3108)

³The syntax of URL is based on *RFC 2396*, and we follow the convention that a URL has four components “<scheme>://<authority><path>? <query>” and we concentrate on the *host* part of an *authority*.

Using this heuristic classification, we found 2233 non-community web sites (out of 18,201) contributing 4156 FOAF documents. We further partitioned the dataset (*GALL*) into seven subsets:

- Groups *G1-G5* for five individual websites contributing over 3,000 URLs:
 - *G1* (www.wasab.dk, 4910 urls) and *G3* (www.kwark.org, 3400 urls) are personal websites mainly for annotating photos
 - *G2* (blog.livedoor.jp, 4266 urls), *G4*(blogs.dion.ne.jp, 3118 urls) are Japanese community websites
 - *G5* ([USERNAME].cocolog-nifty.com ,3108 urls) is a Japanese blog websites.
- Group *GC* contains urls from websites being identified as community,
- Group *GNC* contains all urls from non-community websites.

3.2 Properties of foaf:Person

Since RDF does not have a mechanism of requiring properties for an instance, instances of foaf:Person may come with various kinds and amounts of information. We observed that only 16 properties with the domain *foaf:Person* have been defined in the original FOAF ontology and 140 more have been proposed by other ontologies according to *Swoogle Ontology Dictionary*. In order to evaluate their utility in practice, we collected statistics about the properties being used to describe instances of *foaf:Person*. We found 546 distinct properties used for at least one Person instance, as shown in Table 4. Only 34 properties were used by more than 1% of the FOAF documents. The remaining properties were rarely used FOAF terms (e.g., foaf:yahooChatID), misspelled terms (e.g. foaf:firstname) or relatively new and experimental terms (e.g. foaf:mailbox and http://purl.org/vocab/relationship/spouseof).

Table 4: Property usage in FOAF dataset

property (in URIref form)	document usage		instance usage	
http://xmlns.com/foaf/0.1/mbox_sha1sum	43561	87.56%	114981	55.44%
http://xmlns.com/foaf/0.1/name	34951	70.25%	121498	58.58%
http://xmlns.com/foaf/0.1/nick	33584	67.51%	88217	42.53%
http://xmlns.com/foaf/0.1/weblog	27575	55.43%	70620	34.05%
http://xmlns.com/foaf/0.1/homepage	18712	37.61%	56398	27.19%
http://www.w3.org/2000/01/rdf-schema#seeAlso	18588	37.36%	102589	49.46%
http://xmlns.com/foaf/0.1/knows	13972	28.08%	14686	7.08%
http://xmlns.com/foaf/0.1/depiction	11340	22.79%	12161	5.86%
http://purl.org/vocab/bio/0.1/olb	9318	18.73%	9320	4.49%
http://xmlns.com/foaf/0.1/img	8706	17.50%	8866	4.27%
http://xmlns.com/foaf/0.1/surname	6576	13.22%	9538	4.60%
http://xmlns.com/foaf/0.1/givenname	6530	13.13%	8162	3.94%
http://xmlns.com/foaf/0.1/mbox	5327	10.71%	8463	4.08%
http://xmlns.com/foaf/0.1/firstName	4051	8.14%	6019	2.90%
http://xmlns.com/foaf/0.1/page	3795	7.63%	3851	1.86%

Figure 1 lists the 15 most frequently used terms in FOAF dataset and the percentage of the documents which use each. We associate two types of property usage with two context: (i) document usage reflects the preference of the authors' own personal information and (ii) instance usage reflects the preference of publishing the referred persons' information. Hence, we may find that *name*, *mbox_sha1sum*, are *rdfs:seeAlso* are preferred to describe a link to an author's friends.

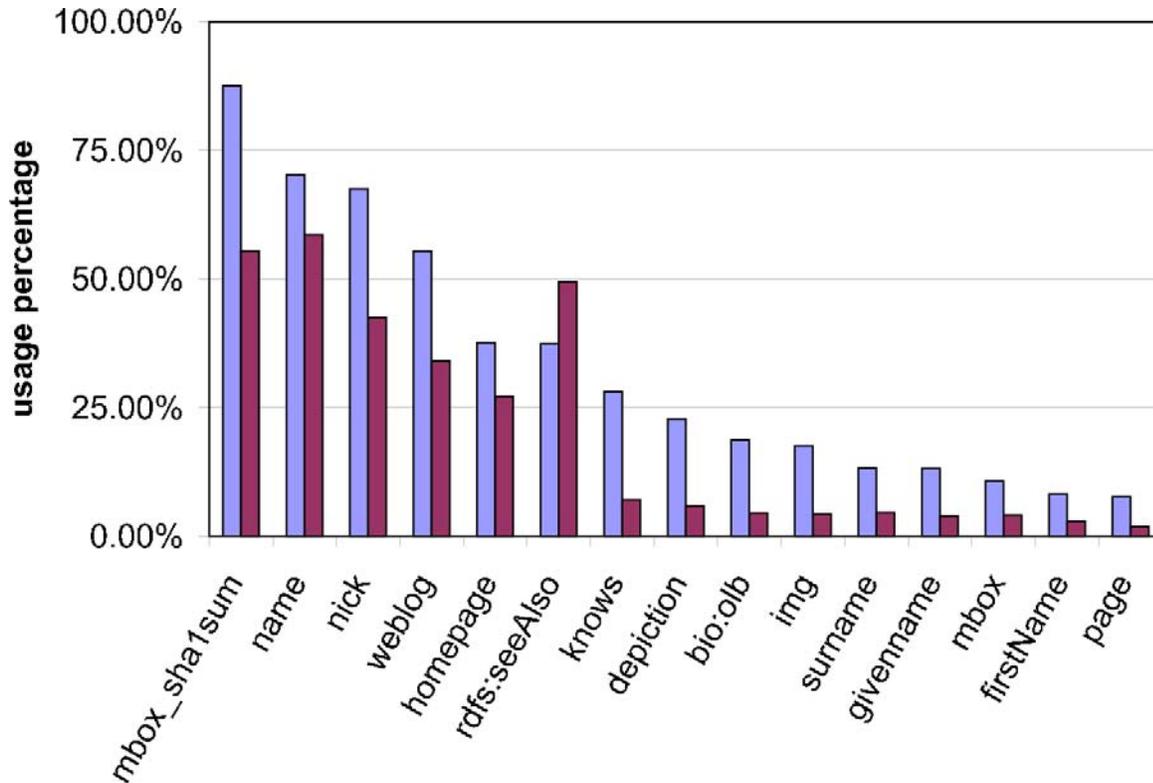


Figure 1: Document/Instance usage of Best used properties

We also observed the impact of community websites in property usage as shown in figure 2. The statistics shows that (i) community websites usually make *mbox_sha1sum*, *weblog* and *nick* mandatory to all their users' profiles, and they may miss some properties, e.g. *depiction* for *G1*, *homepage* for *G2* and *surname* for *G3*; (ii) non-community website authors prefer *name* to *mbox_sha1sum*, and they usually publish their *homepage*, *mbox*, *depiction* (personal photographs), *first name*, *surname*, and *friends*; and (iii) the vocabulary used by community websites are limited in small size (*G1*:58, *G2*:8, *G3*:48, *G4*:8, and *G5*:10) while non-community contribute a very large vocabulary (522 properties). These facts indicate that community websites could skew the overall statistics of FOAF dataset through to their large amount of data; hence identifying community websites is critical to a fair evaluation on the popularity of person property.

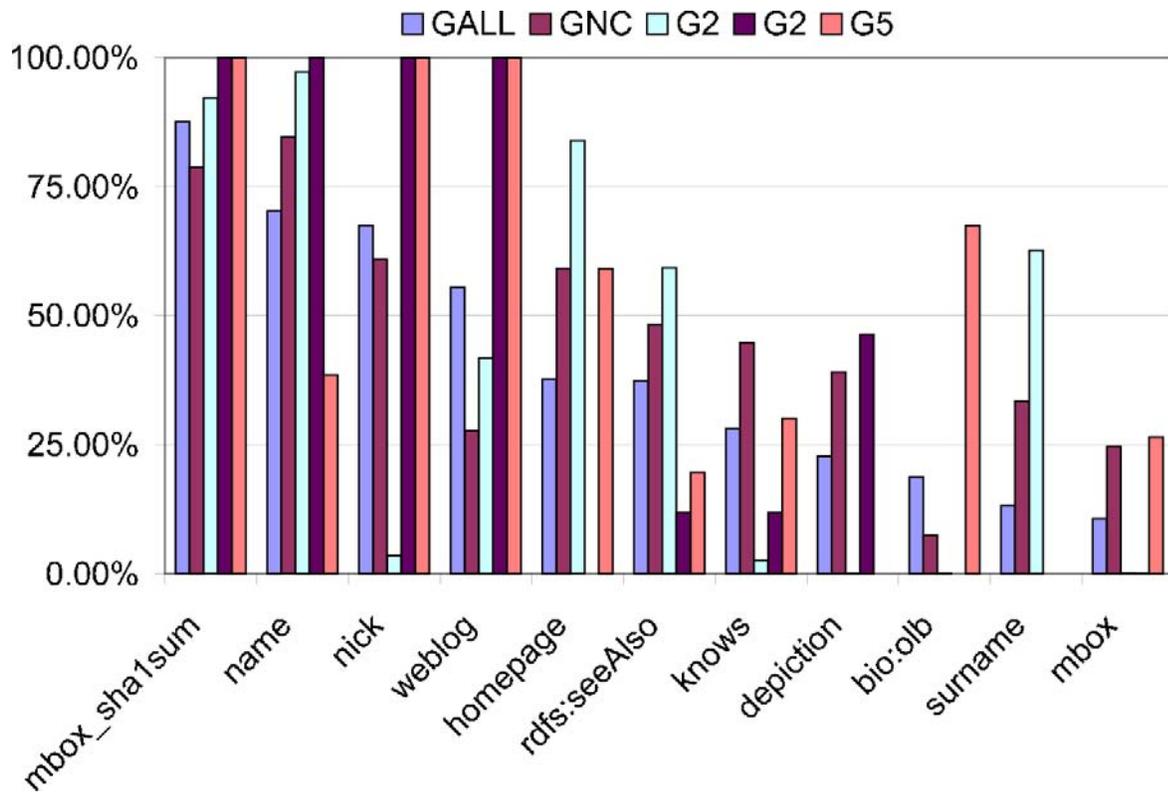


Figure 2: Document usage of best used properties in different groups

3.3 Creators and Referred Persons

All FOAF documents, whether manually or automatically generated, usually require that a person provide the data. Besides the creators' personal information, other persons' information are typically mentioned even when they have not published their own FOAF profile. For example, the email of Dr. *Benjamin Groszof*, a MIT professor, is reported by a document in our FOAF dataset even though he has not published any FOAF document himself.

We classify the person instances into two categories: the **creators** who input their personal profiles and maintain FOAF homepages, and the **referred persons** who are only mentioned by the creators. To this end, we adopt a simple heuristic: the referred persons usually have relatively small amount of triples while the creators have much more. As show in figure 3, we select seven as threshold since there is a sharp drop between seven and eight; and we result in 21,843 (10.53%) creators and 185,570 (89.47%). Another heuristic to identify the creator is to find the one person instance which is not the object of a *foaf:knows* relation.

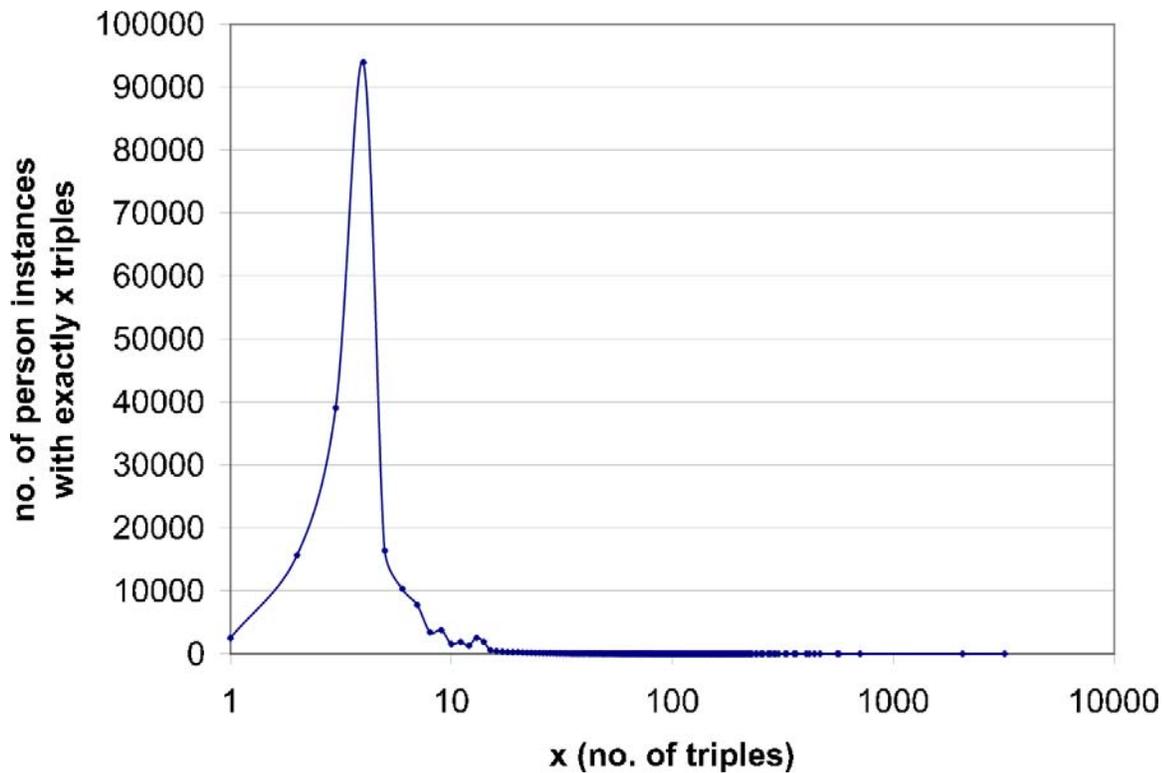


Figure 3: Distribution of no. of triples per person instance

4 Analyzing FOAF social network

We briefly outline two applications involving FOAF data. The first involves the integration and fusion of information associated with individuals. This makes good use of the fact that some FOAF properties can be declared as “inverse functional” and thus offer evidence that two individual FOAF Person nodes describe the same person. The second use is to use FOAF data collected from the Web as data about large scale social networks. FOAF data describing millions of people can readily be collected from the Web today offering new opportunities to explore and test social networking tools, theories and applications.

4.1 Fusing Distributed Personal Information

One of the principles of the Semantic Web is that “anyone is allowed to say anything about any resource”. For example, document *D1* can make assertions about individuals introduced in document *D2*. Since FOAF is based on RDF, this allows one person to assert information about others, be they friends, acquaintances or complete strangers. Hence information about an individual may be spread across a number of FOAF documents in a collection, providing a kind of community view that mirrors the person’s view in the community of people. When a person is described in more than one FOAF documents, we must fuse information from multiple sources and generate aggregated information about the person.

4.2 Person Identifiers

In FOAF data, two foaf:Person instances can be identified as describing the same person in one of two ways. The first is through by URI: two non-anonymous individuals sharing the same URIref in RDF graph can be fused. The second is via assertions involving an OWL InverseFunctional-Property. The FOAF ontology semantics defines unique identifiers of person, such as *foaf:mbox*, *foaf:mbox_sha1sum*, *foaf:homepage* and *foaf:weblog*, which are ideal clues to information fusion. In our FOAF dataset we found 644 URIrefs, 11,405 mbox_sha1sums, 6,099 homepages, 3563 weblogs, and 757 mboxes being used as the identifiers of at least two person instances.

4.3 Fusing person information

Figure 4 demonstrates the result of fusing Dr. Tim Finin’s personal information from 12 sources. We found two different values of *foaf:name* from two different sources in this case: (i) *Tim Finin* as stated by his FOAF profile and (ii) *Timothy W. Finin*” as mentioned in <http://www-2.cs.cmu.edu/People/fgandon/foaf.rdf>. The latter is in fact the unique author identifier in DBLP⁴.

DEMO3: Fuse FOAF Person Information [FLINK](#) [GOOGLE](#) [YAHOO](#) [HOME](#)



foaf:Tim Finin

FOAF provenance	
1	http://www.cs.umbc.edu/~hchen4/harrychen.n3
2	http://www.cs.umbc.edu/~hchen4/harrychen.rdf
3	http://www.cs.umbc.edu/~finin/finin.rdf
4	http://www.csee.umbc.edu/~dingli1/foaf.rdf
5	http://www.cs.umbc.edu/~hchen4/foaf.rdf
6	http://www.cs.umbc.edu/~finin/foaf.rdf
7	http://www.csee.umbc.edu/~finin/foaf.rdf
8	http://www.csee.umbc.edu/~finin/finin.rdf
9	http://www-2.cs.cmu.edu/People/fgandon/foaf.rdf
10	http://www.cs.umbc.edu/~kolari1/kolari-foaf.rdf
11	http://www.cs.umbc.edu/~finin/finin.rdf
12	http://www.cs.umbc.edu/~dingli1/foaf.rdf
13	http://lsdis.cs.uga.edu/~amit/foaf.rdf
14	http://trust.mindswap.org/trustFiles/157.owl

Details about this person		no. of source docs
rdfs:seeAlso	http://www.cs.umbc.edu/~finin/finin.rdf	3
rdfs:seeAlso	http://www.cs.umbc.edu/~finin/foaf.rdf	2
rdfs:seeAlso	http://umbc.edu/~finin/foaf.rdf	1
foaf:aimChatID	timFinin	5
foaf:birthDate	1949-08-04	5
foaf:depiction	http://umbc.edu/~finin/passport.gif	5
foaf:firstName	Tim	5
foaf:homepage	http://umbc.edu/~finin/	5
foaf:mbox	mailto:finin@cs.umbc.edu	7
foaf:mbox	mailto:finin@csee.umbc.edu	5
foaf:mbox	mailto:finin@umbc.edu	5
foaf:mbox_sha1sum	9da08e2b4dc670d9254ab4a4b4d61637fed3b18f	9
foaf:mbox_sha1sum	49953f47b9c33484a753eaf14102af56c0148d37	8
foaf:mbox_sha1sum	8b4d969b2d7dbe0fe5bfc4e069cc2c8a33cf16f4	5
foaf:myersBriggs	ENTP	5
foaf:name	Tim Finin	12
foaf:name	Timothy W. Finin	1
foaf:nick	Tim	5
foaf:phone	tel:+1-410-455-3522	5
foaf:plan	http://www.cs.umbc.edu/~finin/schedule.html	5
foaf:publications	http://www.cs.umbc.edu/%7Efinin/cv/index.shtml#publications	5
foaf:schoolHomepage	http://web.mit.edu/	5
foaf:surname	Finin	5
foaf:weblog	http://ebiquity.umbc.edu/v2.1/blogger/	5
foaf:workInfoHomepage	http://umbc.edu/~finin/	5
foaf:workplaceHomepage	http://umbc.edu/	5

Figure 4: Fusing Dr. Tim Finin’s personal information

Caution should be taken in merging information from multiple FOAF documents since some of the facts may be wrong and the collection of facts may contain contradictions. Small errors in FOAF documents can lead to unexpected results. For example, some FOAF documents from

⁴<http://www.informatik.uni-trier.de/ley/db/>

blog.livedoor.jp, e.g. <http://blog.livedoor.jp/rusa95/foaf00756.rdf>, mistakenly assign the same `mbox_sha1sum` to different people from 4835 FOAF documents. We also found that Dr. Jim Hendler is wrongly fused with Norman Walsh by a FOAF document in which `foaf:mbox_sha1sum` was mistakenly associated with Norman's email-hash.

4.4 Social network analysis

A collection of distributed FOAF documents may constitute a social network. The `foaf:knows` relation can link one individual of `foaf:Person` to another. The FOAF dataset contains 131,314 triples and produces 109,470 `foaf:knows` relations among 49861 instances in FOAF dataset after fusing person. We focused on a smaller portion of that big social network – the emerging social networks in the distributed Semantic Web, which is different than those emerged from a centralized community website. Therefore, our analyses only concerned FOAF documents from non-blog websites. We found many instances followed Zipf's distributions [**Error! Reference source not found.**], so all the figures of distribution were plotted on log-log scale.

4.4.1 Social Network from dataset GNC

We selected about 4156 FOAF documents containing 32727 FOAF person instances before fusing persons. After fusing, we obtain a social network SN_{GNC} with 15,630 `foaf:knows` relations among 26,788 persons. Only 2799 (10%) persons are really fused from at least two original person instances. People fused from many sources could be either social authorities, who are known by many people, or Semantic Web experts (blogger as well), who maintain a fairly large amount of FOAF documents. The top 10 people are list as the following with the amount of original instances they fused from:

- **social authorities**, who are known by many people. For example, Danny Ayers(386), Dan Brickley (199), Libby Miller (133), Edd Dumbill (76), Morten Frederiksen (48), Charles McCathieNevile (39), Dan Connolly (35), Marc Canter (33), Peter Mika(32).
- **Semantic Web experts**, who are usually an active blogger and maintain a “personal” website with large amount of FOAF documents. For example, Christoph Görn (719), Ian Davis (360), Christopher Schmidt(196), Jim Ley(124), Vincent Tabard (71), Masahide Kan-zaki(60)
- **figure in photo**, who have been mentioned by a lot of photo annotation. ONO Hiroki (134), Libby Miller (133), Gregory Todd Williams (61)

4.4.2 Patterns of Degree

Degree analysis is an important tool in social network analysis. Our analyses were based on 15,630 ‘knows’ links within GNC. Figure 6 and 5 shows the distributions of in-degrees and out-degrees respectively. It is shown that only a few fused persons have more than one in-degrees or out-degrees. In fact, among the 26,788 fused persons, only 11.62% of them have both in-links and out-links, and 78.11% of them have only one in-link. All this statistics indicates the sparseness of the SN_{GNC} .

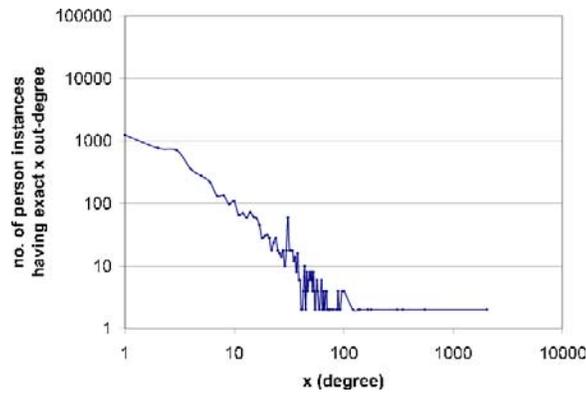


Figure 5: Out-degree distribution per group

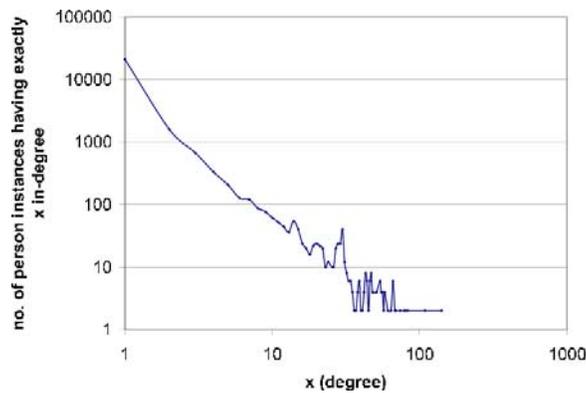


Figure 6: In-degree distribution per group

4.4.3 Patterns of Connected Components

There are 842 components in SN_{GNC} with average size 16. The distribution of component size is highly skewed as shown in figure **Error! Reference source not found.**: there is one very large component with 7,111 fused individuals and the second with only 549 (less than 10% of the size of the first). We note that the large component was fused due to errors in the FOAF documents, which mistakenly assigned the same *foaf:mbox_sha1sum* to many different individuals.

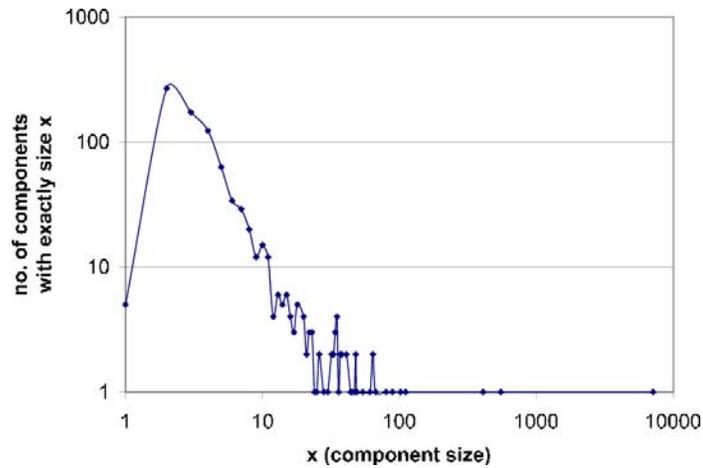


Figure 7: Distribution of component size

The inherent nature of FOAF publishing makes the star-shaped component shown in Figure 8⁵ common. It typically arises when an individual publishes a FOAF document describing a set of people with whom they have a *foaf:knows* relationship.

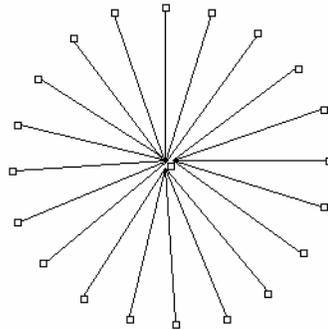
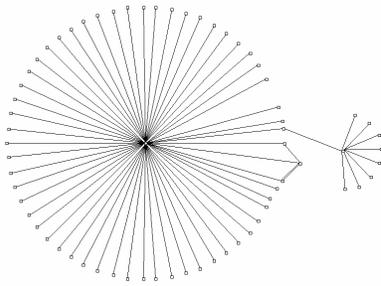


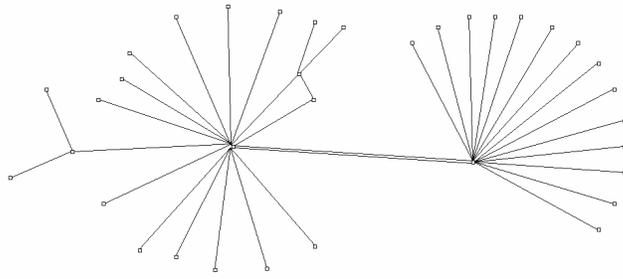
Figure 8: Star shaped components are common in naturally occurring FOAF profiles.

As more people publish FOAF profiles, the star configurations of the early adopters may get their influence spread in bud mode as depicted in figure 9a or hook up with each other though bi-directional bridges, as shown in figure 9b.

⁵Figure **Error! Reference source not found.** - 11 were generated by the “Otter” network visualization tool [**Error! Reference source not found.**].



(9a) expanding the next level



(9b) two social networks being bridged

Figure 9: Component growth models

The second largest component in SN_{GNC} , as shown in figure 10 with 546 nodes and 771 directed edges, turns out to be a proof of the above social network growth models. It features several hubs with very high out-degree, plus several other nodes with extremely high betweenness value staying between those hub nodes.

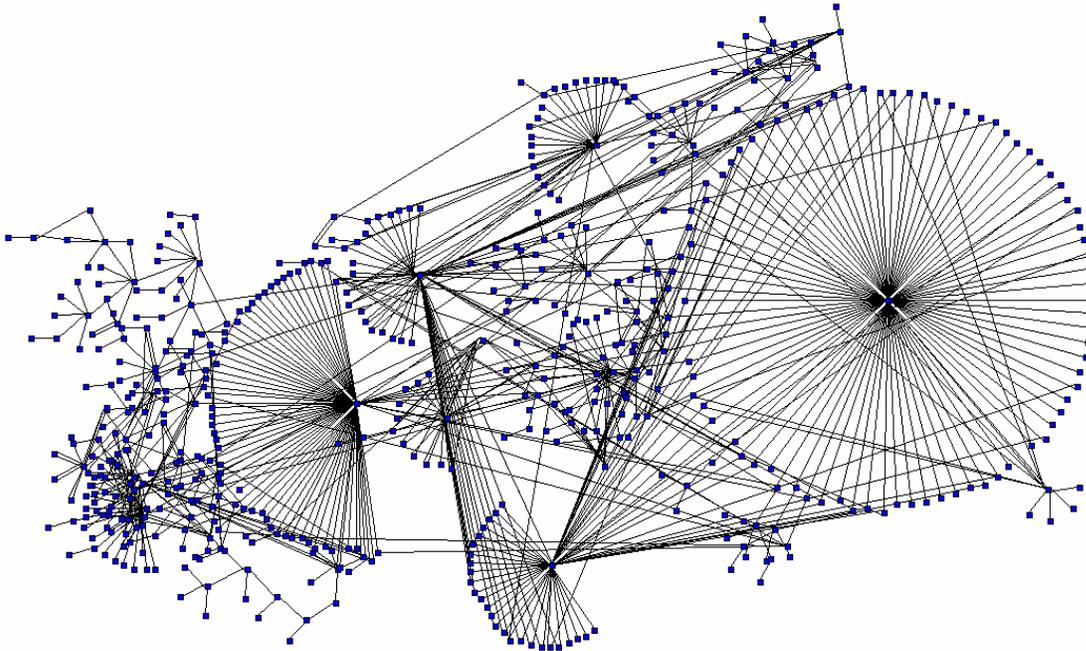


Figure 10: The second largest component in SN_{GNC}

5 Conclusions

The Semantic Web offers an ambitious vision of an Internet populated with intelligent agents and services able to exchange information, tasks and knowledge using simple protocols coupled with a rich knowledge representation language. Exploring the roadmap leading toward this vision will take some time. The Semantic Web languages RDF and OWL are a promising beginning. One of

the first wide-spread applications of RDF is the representation of social networks – individuals, their properties and the relationships among them. The current interest in social networks and the immediate applications to online virtual communities have made the FOAF ontology the widely used on the web at this writing. Studying how FOAF is being used provides a good test case for the larger questions and issues involving the adoption of Semantic Web concepts and technologies.

We presented a novel perspective of the Semantic Web by linking machine-readable descriptions of people, i.e. FOAF documents, with published personal relationships. This complements the ontology-based view of the Semantic Web. We also proposed a heuristic approach to identifying and discovering FOAF documents from the Web and extracting information about people from these FOAF documents. This approach provides a means of transcending the boundaries of individual FOAF documents, fusing information about a person from multiple documents. The analysis of FOAF network pattern also lent itself to unique social network structures in the Semantic Web.

FOAF networks provide a snapshot of the FOAF user community encoded in the constituent *foaf:knows* relations. More importantly, connection patterns among FOAF documents offer a persons orientation to the conventional Web of HTML documents. The visualization of highly connected FOAF networks is informative and revealing. As the number of FOAF users grows, the approach presented in this paper can be used to discover existing and emerging online communities.

Acknowledgements

Partial support for this research was provided by DARPA contract F30602-00-0591 and by NSF awards NSF-ITR-IIS-0326460 and NSF-ITR-IDM-0219649.

References

- [**dc**] The Dublin core element set v1.1 namespace providing access to its content by means of an RDF schema. <http://purl.org/dc/elements/1.1/>.
- [**foaf**] Foaf vocabulary specification version 1.0. <http://xmlns.com/foaf/0.1/>.
- [**foaf-project**] The friend of a friend (foaf) project. <http://www.foaf-project.org/>.
- [**rss**] Rdf site summary. <http://web.resource.org/rss/1.0/spec>.
- Adamic, L. A., Buyukkokten, O. Adar, E. 2003, June. A social network caught in the web. First Monday, 8(6).
- Chen, C. (1999). visualising semantic spaces and author co-citation networks in digital libraries. Inf. Process. Manage., 35(3), 401–420.
- Ding, L., Zhou, L., Finin, T. Joshi, A. (2005). How the Semantic Web is Being Used: An Analysis of FOAF, Proceedings of the 38th International Conference on System Sciences, Hawaii, January 2005.

- Ding, L., Finin, T., Joshi, A., Pan, R., Cost, R. S., Peng, Y., Reddivari, P., Doshi, V. C., Sachs, J. (2004). Swoogle: A search and metadata engine for the Semantic Web. Proceedings of the thirteenth ACM conference on information and knowledge management.
- Ding, L., Kolari, P., Ganjugunte, S., Finin, T., Joshi, A. (2004). Modeling and evaluating trust network inference. Seventh international workshop on trust in agent societies, The Third International Joint Conference on Autonomous Agents and Multi Agent Systems, New York, July 2004.
- Ding, L., Zhou, L. Finin, T. (2003). Trust based knowledge outsourcing for Semantic Web agents. Proceedings of IEEE/WIC international conference on web intelligence.
- Dumbill, E. (2002a). Finding friends with XML and RDF. IBM's XML Watch, <http://www-106.ibm.com/developerworks/xml/library/x-foaf.html>.
- Dumbill, E. (2002b). Support online communities with foaf: How the friend-of-a-friend vocabulary addresses issues of accountability and privacy. IBM's XML Watch, <http://www-106.ibm.com/developerworks/xml/library/x-foaf2.html>.
- Dumbill, E. (2003). Tracking provenance of RDF data. IBM's XML Watch, <http://www-106.ibm.com/developerworks/xml/library/x-rdfprov.html>.
- Garton, L., Haythornthwaite, C. Haythornthwaite, C. (1997). Studying online social networks. Journal of Computer-Mediated Communication, 3.
- Golbeck, J., Parsia, B. Hendler, J. (2003). Trust networks on the Semantic Web. Proceedings of cooperative intelligent agents.
- Grimnes, G. A., Edwards, P. Preece, A. (2004). Learning meta-descriptions of the FOAF network. Proceedings of international Semantic Web conference.
- Hayes, P. (Eds.). (2004). Rdf semantics (w3c recommendation, 10 February 2004). <http://www.w3.org/TR/2004/REC-rdf-mt-20040210/>.
- Huffaker, B. (1998). Otter: A general-purpose network visualization tool.
- Klyne, G. Carroll, J. J. (Eds.). (2004). Resource description framework (RDF): Concepts and abstract syntax. <http://www.w3.org/TR/2004/REC-rdf-concepts-20040210/>.
- Kumar, R., Raghavan, P., Rajagopalan, S. Tomkins, A. 1999. Trawling the web for emerging cyber-communities. Proceeding of the eighth international conference on world wide web (1481–1493). Elsevier North-Holland, Inc.
- Milgram, S. (1967). The small world problem. Psychology Today, 1(1), 60–67.
- Schreiber, G. Dean, M. (). 2004, February. Owl web ontology language reference. <http://www.w3.org/TR/2004/REC-owl-ref-20040210/>.
- Xu, J. Chen, H. (2003). Untangling criminal networks: A case study. Intelligence and security informatics, first NSF/NIJ symposium (2665, 232-248).
- Zipf, G. K. (1932). Selected studies of the principle of relative frequency in language. Harvard University Press.