Project no. 34721

TAGora

Semiotic Dynamics in Online Social Communities

http://www.tagora-project.eu

Sixth Framework Programme (FP6)

Future and Emerging Technologies of the Information Society Technologies (IST-FET Priority)

D2.3 Interim Report on tagging systems update and usage

Period covered: from 01/06/2007 to 31/05/2008
Start date of project: June 1st, 2006
Due date of deliverable: July 15th, 2008
Distribution: Public

Date of preparation: 31/05/2008
Duration: 36 months
Actual submission date: June 20th, 2008
Status: Final

Project coordinator: Vittorio Loreto
Project coordinator organisation name: Kassel
Lead contractor for these deliverables: Kassel

Executive Summary

This deliverable provides an interim report about the development of the following systems and draws conclusions for the work in the third project year:

BibSonomy – a social resource sharing system for bookmarks and publications

Tagster – Folksonomy Peer-to-Peer System for Sharing Multimedia Data

zexe.net – A Community Memory for Representing Daily Experiences using Folksonomies

Ikoru – A Test-bed for Collaborative Tagging and Content-Based Analysis

TAGnet – a tool for awareness and management of personal metadata

MyTag – a tool for integrate folksonomies

The aim of these systems in Tagora is twofold: they shall support the collection of real-world user data for experiments, and provide a platform for experimenting how different kinds of user interaction influence the evolution of the resulting data over time.

The first system, *BibSonomy*, allows users collaborative organizing and sharing of bookmark collections and publication lists. A basic version of BibSonomy has been online before the start of the project. Within Tagora, we have extended its functionality to attract a significant number of users, and have provided means for a systematic generation of data for experiments.

Tagster is a system for collaboratively organizing and sharing multimedia data in a peer-to-peer network. It is completely decentralized and provides the same functionalities as common centralized folksonomy systems like Flickr, delicious or BibSonomy. Besides the basic tagging and browsing support it is also used as a test bed for implementing new features like distributed management of metadata statistics. The tagging data gathered in the system will following be used for further analysis and comparison with the dataset from centralized tagging systems.

*canal*MOTOBOY* and *GENEVE*accessible* are the latest projects of the zexe.net initiative. Both of these projects make an intensive use of tagging. In *canal*MOTOBOY*, 15 motorcycle messengers in Sao Paulo, Brazil, use multimedia mobile phones to capture images and videos of their daily life. They use tags to describe these contents, which they publish on the web. The *GENEVE*accessible* project involves handicapped people in Geneva, Switzerland. They use multimedia phones equipped with GPS to create maps of their city's accessibility. They use tags to describe the images of obstacles they find in their way. By publishing these tagged and geo-referenced images on the web, they effectively build an intelligent, collaborative map which is immediately available to the public. These projects have enabled members of TAGora to study the dynamics of tagging in small-scale groups with shared interests. In particular, these projects provide two contrasting examples. While *canal*MOTOBOY* is totally open-ended, *GENEVE*accessible* has very specific goals. We have studied how these different scopes affect the projects' folksonomies.

The *Ikoru* system, developed at Sony CSL, is primarily used to experiment with collaborative tagging and content-based analysis. The project consists of a server-side component and a Web interface, which can be viewed at `http://www.ikoru.net`.

The *TAGnet* system is a prototype (not yet available on the web) designed to provide users with a reflexive tool to expose regularities and patterns in their own tag-based annotations. Tagging patterns can reveal a lot about a user's experience, her interests and her emergent conceptualizations, but users are not aware of these patterns until these regularities are made explicit by means of data analysis and state-of-the-art visualization. TAGnet currently focuses on Flickr users, providing them with a "semantic mirror". It is conceived as a web application that provides users with actionable meaning on their own metadata. In perspective it will be also used as a tool to explore

emergent conceptualizations and tag ranking strategies. This application was not initially foreseen, and was set up to exploit the results of WP3 and WP4 on the structure of tag co-occurrence networks.

The *MyTag* system aims at solving the limitations of current tagging platforms by enabling cross-media search across images, video, and social bookmarks. It offers transparent access to different single-media platforms currently including Flickr, YouTube, and del.icio.us. The search function can be personalized in two directions. First, MyTag users can restrict the search to the resources uploaded by the user him/herself. Second, the website uses an implicit user feedback mechanism to personalize the output of a query to MyTag by ranking results according to the user's personomy. The personomy is built without additional effort by the previous queries entered by the user, in contrast with other tagging platforms, such as Flickr or del.icio.us, where an explicit feedback is required in order to personalize the ranking of the results presented to the user.

Contents

List of Figures

Chapter 1

BibSonomy — A Social Resource Sharing System for Bookmarks and Publications

This chapter describes the extended functionality of BibSonomy that we have implemented during the project in order to increase user activities. This strategy has been successful, as indicated by Figure 1.1, which shows the increase of the number of users since the the last report in May 2007. New features are announced on a weekly basis on `http://bibsonomy.blogspot.com/`.

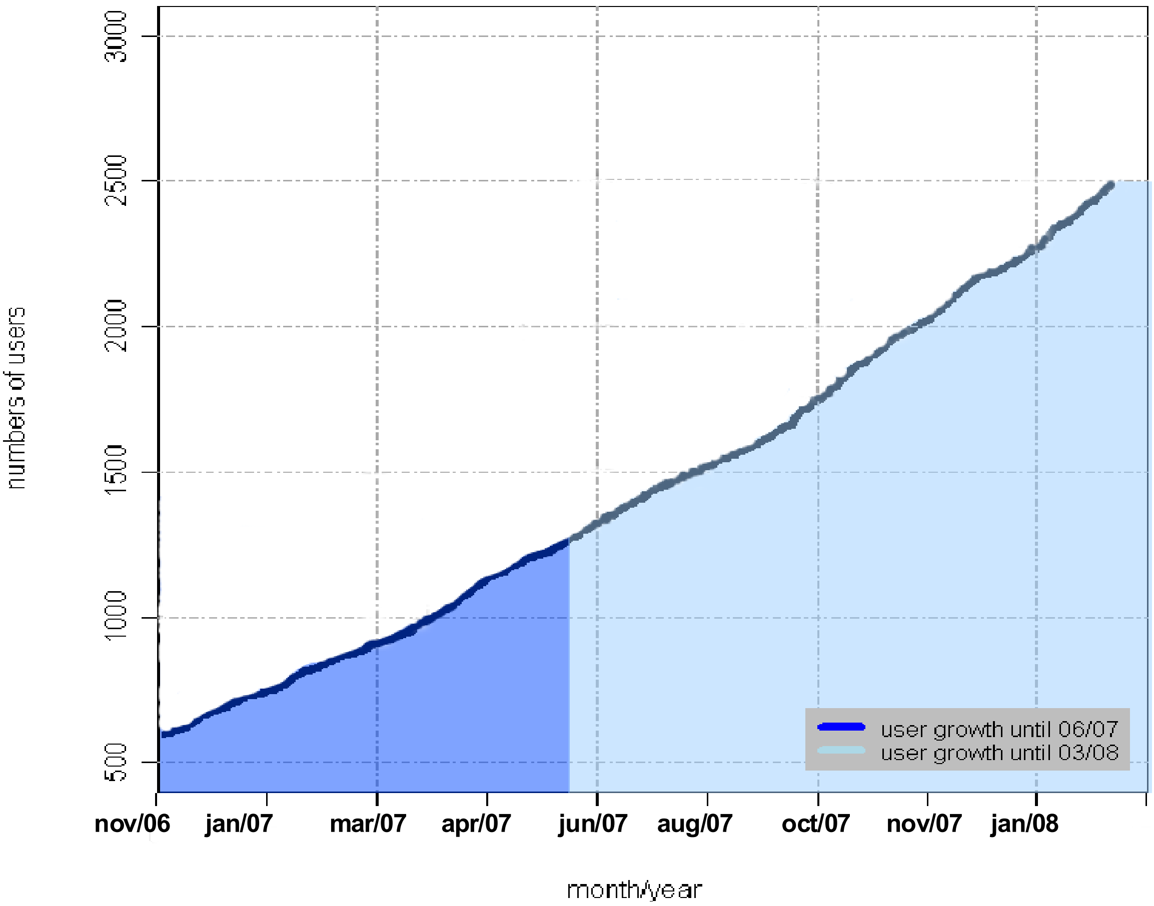

Figure 1.1: Growth of BibSonomy since May 2007

1.1 Supporting community building in BibSonomy

1.1.1 Explicitly defined communities/groups

Currently, BibSonomy contains overall 61 groups, which actively sharing and collecting resources. The groups are mostly research groups or participants of European projects. To facilitate the idea of group working, the system offers a common tag cloud for each group, which is characterized by a filter to restrict the view on the tags. In this way the sharing of resources is restricted to a group. To become a group member, the users have to apply for the group of interest. The group administrator will be informed by email and can accept or reject the request. After a user has joint a group, he can see all entries of this group (i. e., also the group internal resources).

1.1.2 Searching enhances implicit community building.

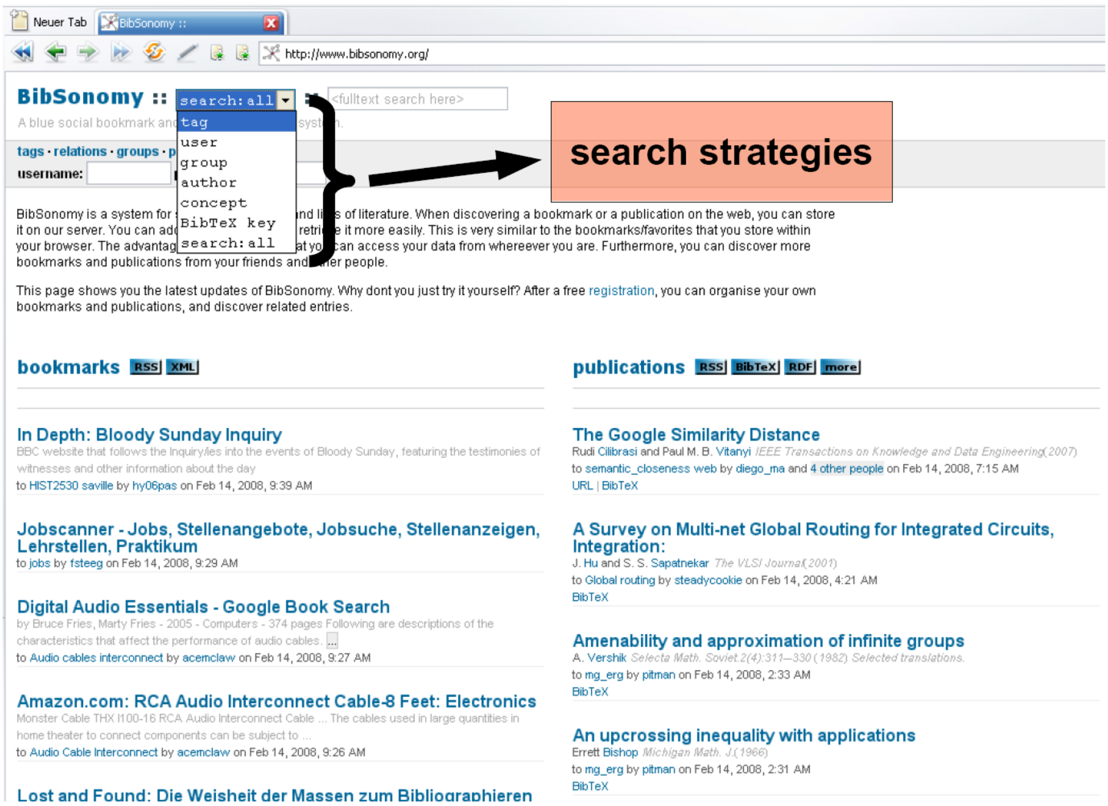

Figure 1.2: Search strategies in Bibsonomy

To find content from users who used similar tags as oneself, BibSonomy supports currently more user flexibility regarding different search mechanisms. The following search strategies were implemented after 1st of June 2007 and are available online for navigation (see Fig. 1.2).

Searching by author BibSonomy allows users to search in publications via author names. The author search is implemented based on the MySQL full text search feature of the MyISAM database engine. The system copies all the author information of a publication into a text field of a MyISAM table, yielding a very fast search functionality.

The simplest way to search for an author is to search for the last name. That is what the system supports right now and is currently called an ‘author page’. The author search results in a list with all publications together with a tag cloud describing the topics of the author based on the tags that are attached to his publications in the system.

Figure 1.3: *SUPERTAG* <- *SUBTAG* relation in Bibsonomy.

Searching by concepts BibSonomy allows users to structure the content via *SUPERTAG* <- *SUBTAG* relations (see Fig. 1.3). When 'concept' is selected in the pull-down menu in Fig. 1.2, then the search does not result only in the resources that are annotated with the given tag, but returns also all resources which are annotated by at least one subtag.

Ranking Another major extension contained in BibSonomy is a mechanism to rank resources and users for a given tag by relevance. To this end, we have implemented the *FolkRank* algorithm (Hotho et al., 2006). Its idea is similar to Google's PageRank algorithm, i.e. it analyzes the link structure between users, tags and resources in order to calculate the relevance (more details in Deliverable 3.2).

1.2 Integration with 3rd party products

We made several steps to combine BibSonomy with third party products.

Zope Zope[1] is an open source application server for building content management systems, intranets, portals, and custom applications. Publication lists, link lists, and tag clouds can be dynamically integrated into Zope web pages, using the *KebasData* product (Zope, 2002).

Wiki -and Webblog Software BibSonomy also supports Wiki -and Webblog Software. To integrate BibSonomy data into an XWiki-Page, one has to install the XWiki RSS Feed Plugin (XWi, 2004). Then, the data can be imported as RSS Feed. To import BibSonomy data to a WordPress blog, the WordPressBlog plug-in [2] has been implemented.

Digital Libraries meet BibSonomy The Library of the University of Cologne [3] was the first 3rd-party organization that incorporated BibSonomy's services: When searching for books and articles, the results can be stored with one mouse click in a personal bibliography collection at BibSonomy.

The Library of the Institute of Information Sciences at the Saarland University, Saarbruecken [4] also integrated BibSonomy into their literature research interface. In addition to the features provided by the KUG library (i.e. the direct posting of search results) links are provided to retrieve further articles from BibSonomy by author name.

The Library of the University of Heidelberg[5] was following this service, and a similar implementation at the University of Kassel[6] is coming up.

Moodle Moodle is a popular e-learning platform for students and lectures (Dougiamas, 1999). BibSonomy can be integrated to enhance course descriptions and e-learning projects by providing the corresponding literature also via RSS-feeds.[7]

[1] http://www.zope.org/
[2] http://www.christianschenk.org/projects/wordpress-bibsonomy-plugin/
[3] http://kug.ub.uni-koeln.de/
[4] http://is.uni-sb.de/vibi/suchen.html
[5] http://katalog.ub.uni-heidelberg.de
[6] http://opac.bibliothek.uni-kassel.de
[7] http://educampus.uni-kassel.de/

1.3 BibSonomy as Web Service

BibSonomy's REST API BibSonomy now has an application programming interface (API) which allows external applications to interact with BibSonomy. It is restful API which provides a simple access to all data of BibSonomy. One example application which uses the API to access BibSonomies data is the stand alone BibTeX Manager JabRef[8]. It is an open source tool to manage bibliographic metadata. We have extended the tool by using the client part of API implementation to establish a connection to BibSonomy. The client is able to store and to retrieve references from BibSonomy. The user tag cloud or the tag cloud of the system is shown after the start and offers an intuitive browsing interface. Tags can be used in the usual way to search for references directly in Jabref, and the retrieved reference list can be imported into Jabref's internal library with one click. This also means that references can be used on a laptop without having network connectivity, and will be synchronized with BibSonomy once reconnected.

For the more technically minded: the API is based on REST.[9] To get all users, the data are requested as "GET /users" over HTTP. To modify a particular user, the expression "PUT /users/<username>" is used, together with an appropriately formatted XML document containing the user data. The API documentation is available on `http://www.bibsonomy.org/help/doc/api.html`.

It is possible to code an application against BibSonomy in about any programming language, although you will have to write all the HTTP and XML wrangling yourself. For the Java language, BibSonomy is also offering a client library (available on `http://www.bibsonomy.org/help/doc/javaclient.html`). We expect that the community will setup client implementations for other programming languages like, python, perl or php, too.

The discussed functionality of BibSonomy as Web Service works within the bounds of proper authorization. To be able to use the API, interested users can obtain an API key by checking the setting page within its user account.

Multi-language support As researchers and students of different nationalities work with BibSonomy, version 2.0 supports multilingual pages. Almost all non-posting pages of BibSonomy are now available in English and German.

1.4 Dissemination of BibSonomy

Mailing List Our intension is to build up a denser network around BibSonomy research. Therefore, we have set up a mailing list *bibsonomy-research@cs.uni-kassel.de*. The list is intended for facilitating exchange about research issues and research around BibSonomy. The update of the BibSonomy dump every three months is also published on the mailing list. Currently, 43 members are participating in the mailing list. Regarding the BibSonomy dump, already 32 people signed up the data agreement for research means.

Cooperation BibSonomy is currently used at the Fraunhofer Institute for Autonomous Intelligent Systems and at SAP Research for internally organizing publications. Primarily, the local implementation serves as first test phase and will be extended after a successful evaluation.

[8] `http://jabref.sourceforge.net/`
[9] `http://en.wikipedia.org/wiki/Representational_State_Transfer`

Conference Support The system contains all accepted papers of the conferences Statphys23,[10] ISWC+ASWC 2007,[11] and ESWC 2008,[12] together with the keywords (tags) that authors have associated with their papers (or that show up in the paper titles). To help conference participants find interesting publications, a web-front-end has been created which shows a tag cloud of the most important keywords. The colour of each tag indicates the track to which most abstracts annotated with that tag belong to. Clicking on a tag (keyword) will retrieve from BibSonomy the abstracts that have been tagged with it. Given the necessary BibTeX entries to store publication abstracts, metadata and associated keywords in BibSonomy, we can provide BibSonomy web front-ends presenting a conference's tag cloud and interests.

Discovery Challenge We are currently organising the Discovery Challenge of the ECML/PKDD 2008 conference.[13] The challenge comprises two tasks: learning tag recommendations, and detecting spam, both based on a BibSonomy dataset. The final results will be published in September 2008. We expect to gain valuable input from the challenge for WP 4. The challenge is organised as Tagora outreach.

Publication The new features of BibSonomy have been presented at the Workshop on Social and Collaborative Construction of Structured Knowledge (CKC 2007) at WWW 2007 (Jäschke et al., 2007).

1.5 Conclusions for the third year

We are going to further extend the functionality of BibSonomy as planned, in order to attract more users and to enlarge our dataset for the experiments. Additionally, we will focus on spam detection and at recommendations, as these have been shown to be crucial features for the usability of the system.

[10] `http://www.bibsonomy.org/events/statphys23`
[11] `http://www.bibsonomy.org/events/iswc2007`
[12] `http://www.bibsonomy.org/events/eswc2008`
[13] `http://www.kde.cs.uni-kassel.de/ws/rsdc08/`

a)

b)

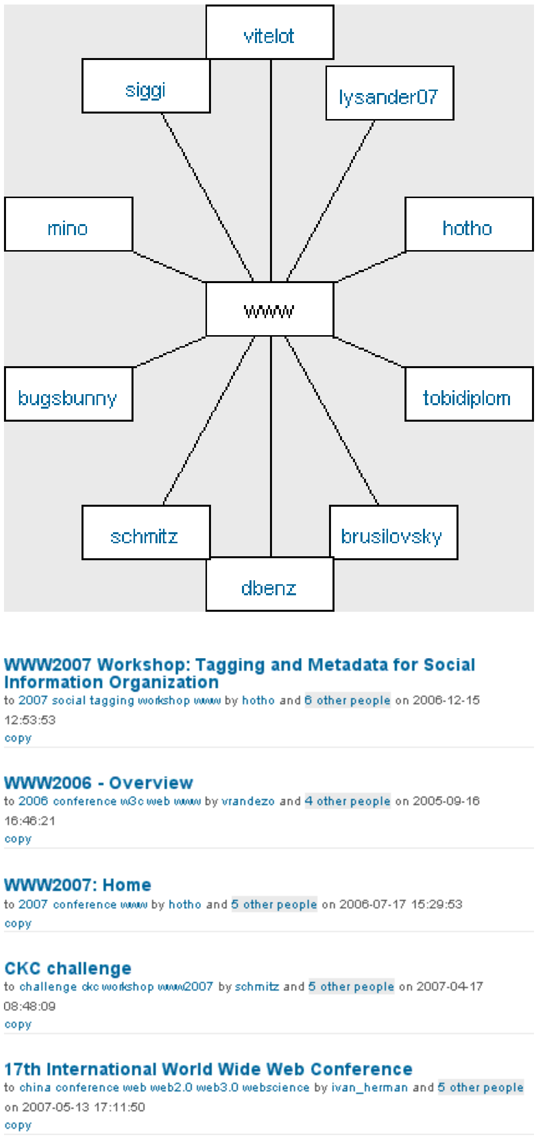

Figure 1.4: a) Ranked users of FolkRank related to the tag 'www'. b) Ranked websites according to this tag.

Chapter 2

Tagster - Folksonomy Peer-to-Peer System for Sharing Multimedia Data

Tagster is a distributed tagging system for sharing multimedia data. It is based on the Semantic Exchange Architecture (SEA) (Franz et al., 2006) which provides the basic functionality for organizing and sharing annotated information resources in a decentralized scenario. Additionally, a mechanism for managing distributed tagging statistics is integrated and the application provides different data views for easily navigating the annotated multimedia data.

2.1 Improvements

Several improvements have been made during the second year of the project. At the end of the first year, the prototype was mainly providing the basic functionality for tagging multimedia data in a decentralized fashion. During the second year we have extended the application in two main directions: the management of distributed tagging statistics and the improvements concerning the user interaction and usability.

Distributed Statistics

To make tagging meta data available to all users in the network, Tagster also puts it in a global index structure. That means each peer in the network stores a fraction of the globally available meta data and the underlying index implementation (we use Bamboo[1]) assures that the stored data can be accessed in a very efficient way. However, the index only stores pairwise relations between users, tags, and resources. Handling more complex information retrieval tasks like similarity computation of users and result ranking would require contacting many peers to gather the necessary information which is apparently highly inefficient. Therefore, we have developed a novel mechanisms for managing distributed statistics, called PINTS (Görlitz et al., 2008).

The basic idea for distributed meta data management is that each peer in the network is maintaining a fraction of the global meta data. With each new tag assignment the responsible index peer updates the respective index information and notifies other peer about the changes if necessary. The similarity computations are based on the cosine similarity of feature vectors, as for example tag clouds. We adapted the well known TF-IDF measure from information retrieval such that each feature combines local and global data like a user's tag frequency and the tag's popularity in the whole network. The problem, however, is to keep the statistics accurate since the propagation of every change of the global index data would cause a high message complexity. Therefore, the PINTS algorithm only propagates data updates if the estimation of the change in the depending

[1] http://bamboo-dht.org

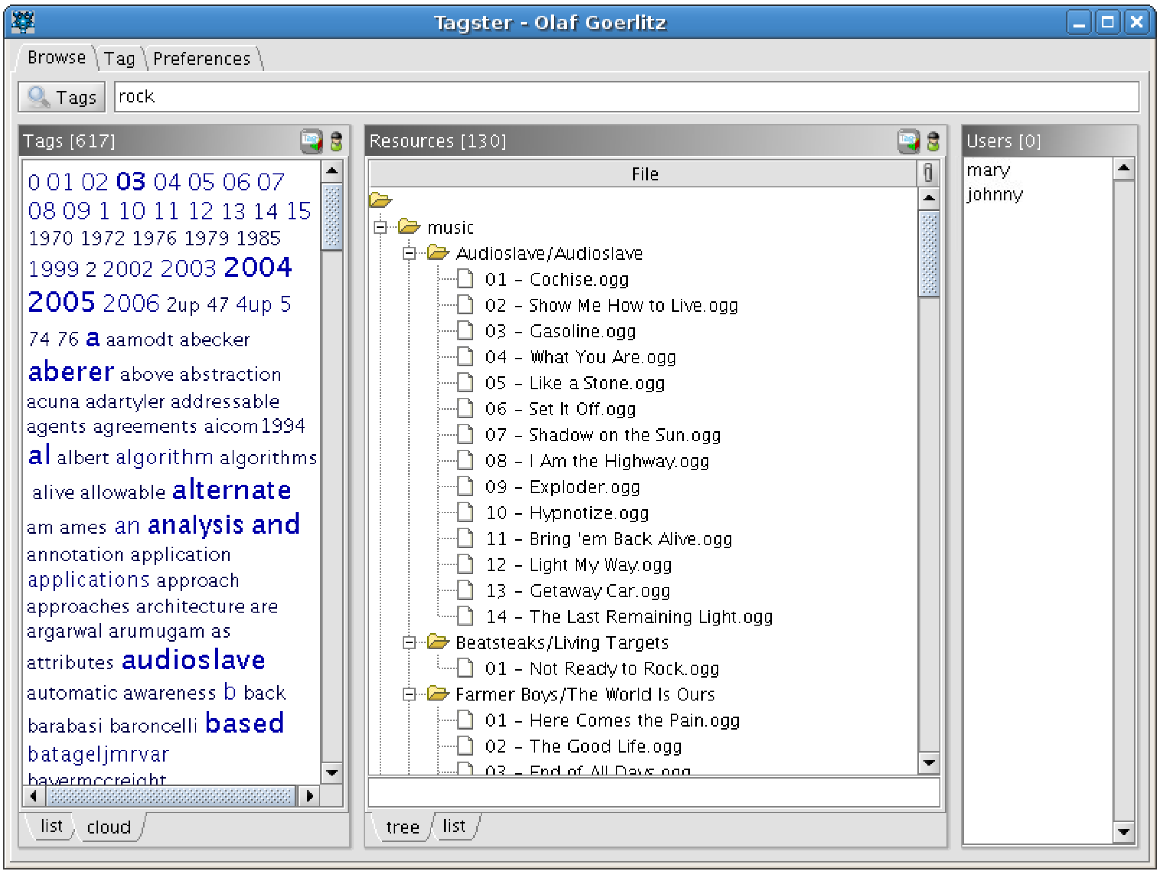

Figure 2.1: Tagster's main view showing the user's tag cloud and all resources tagged with 'rock'.

statistics is higher than a certain threshold. That allows us to maintain accurate distributed statistical information while keeping the message complexity low in the network. PINTS is implemented in Tagster and used to display statistical information like tag clouds.

User Interface

The design of the user interface plays a major role for the usability of a software. In the case of Tagster it is important to have an interface that provides the same functionalities like the centralized folksonomy systems but also includes an intuitive way of navigating though both the personal data on the local machine and the information retrieved from the network.

However, the application's appearance is not the only aspect we consider for a good user interface but also the ease of use, i.e. the simplicity of configuration and setup/joining of the peer-to-peer network.

Navigation elements The adaption of typical navigation elements like tag clouds from the centralized systems is strongly motivated by the fact that user of such systems are already very accustomed to that type of navigation support. Therefore, one goal is to integrate the same or similar data representations such that users can get familiar with Tagster really quick. This includes, for example, the display of related information for the currently selected data items and contextual tag clouds which are a very typical navigation element.

Additionally, we have integrated resource-specific type views, i.e. the user can browse the resources by their associated Mime type. Thus, it is possible to filter a search results such that only images or documents are displayed.

Tagster's local resource organizations allows the user to tag any file on the local harddisk. Thus, also a file's path information is preserved and displayed in the resource view. To better visualize the local resources we have implemented a tree view that orders all local file in their actual folder hierarchy. Resources from the network are displayed without a hierarchy since that information is not returned for privacy reasons.

Additional functionality Since browsing of resources in the network is not enough, we have also implemented a download protocol for directly retrieving files from other users. To download a resource from the network the user just has to click on the download button next to the resources displayed tagging data. Then the file will be retrieved from the owner and saved in the local download folder. All tags already assigned to the resource will be automatically applied to the downloaded file, too.

Tagster's resource view only diplays the typical file information like name and path. To actually see the content of the files we integrated a function to start the appropriate external application that is associated with the respective file. The intention is that the user does not need to switch manually to another application to view his resources. Currently, this function is supported on Windows and Gnome-based Linux systems.

Configuration and Network Setup

Tagster requires some complex configuration settings due to its distributed nature. Especially, the network setup and peer lookup requires specific port settings etc. that the user ideally should not be bothered with. Therefore, we tried to simplify the configuration and setup process as much as possible. The user is only obliged to enter some personal information at the first start of Tagster. All other system settings are initialized to default values and only need to be changed by the user in some rare cases.

Joining a network is the most complicated part of the initialization process since some bootstrapping information is required. To simplify that process we are using a central user registry that keeps track of all active users in the network. During the network initialization phase, which is completely transparent for the user, the Tagster client first retrieves a list of peers that are used for the network rendezvouz. Afterwards, the client sends its own peer information to the registry.

One important problem of the decentralized approach is the uptime of users. Although data replication is used not all of the distributed index information may be available all the time if only a small number of users is online. To compensate this effect we have set up one peer on a server to run 24/7.

Data Gathering

For further analysis we are gathering the tagging meta data of all users in the network and merge it into one large data set. However, gathering the data is not a trivial task as not all users are online at the same time. Although the tagging meta data is completely available in the global index it is not a good idea to copy it from there. This would involve a high message overhead in the network to query all indexed data items. Moreover, the index stores no timestamp information for tag assignments.

Currently, we are relying on the cooperation of the initial test user group to extract the desired data directly from their local repository. However, in the future it will be necessary to have a more sophisticated data gathering mechanism that can incrementally retrieve the data from the online peers in the network without even bothering them.

In the collected dataset we store for each tag assignment the timestamp, the user's ID, the tagged resource's ID, and all assigned tags. Addionally, we log all data queries resulting from the users search and browsing activity. The query logs contain a timestamp, the type of the query (i.e. query for tag, user, or resource) and the queried data item. All gathered data is available at http://isweb.uni-koblenz.de/DataSets/TagsterDataSet and will be regularly updated in the future.

2.2 Installing and using Tagster

The Tagster prototype can be downloaded at http://isweb.uni-koblenz.de/Research/tagster. New software versions will be published when available. The installation is very simple as the application comes as an executable Java archive. The initial setup is guided by a setup dialog and Tagster automatically joins the peer-to-peer network.

Currently, we aim to attract more users, especially from the scientific community. Tagster was first made available to the member of the ISWeb research group in Koblenz. The next step is to distribute it among all partners in the project and their respective research groups. We hope to attract at least 50 user until the end with the majority willing to use Tagster on a regular basis.

2.3 Experiences

We have experienced a strong interest in Tagster. Many people like the idea of sharing their resources in a peer-to-peer fashion, especially with colleagues. However, the attraction of a sufficiently large user group with a long-term interest to use the software is crucial. A user's client should be running for several hours each day. Otherwise, the interest will quickly diminish because nobody can find other peoples resources.

Another severe problem are firewalls. The libraries integrated in Tagster to maintain the distributed index are all coming form research work at universities and unfortunately they are generally lacking a good support for firewall handling. An extension into that direction requires some significant effort since the firewall tunneling techniques are rather complicated and the integration in the libraries is hampered by poor documentation.

2.4 Road map for the third year

The main task at the beginning of year three is the improvement of Tagster's networking capabilities. That is especially the by-passing of firewalls but also the optimization of the availability of the meta data in the network by caching or the use of super-peers. Considering the distributed statistics, we want to extend our approach towards the distributed processing of matrix operations which are required for generating more complex statistics. Finally, the data gathering mechanism will be improved in a way that a continuous collections of tagging data from all peers in the network is possible.

Chapter 3

zexe.net - A Community Memory for Representing Daily Experiences using Folksonomies

The zexe.net system consists of a set of online applications and tools that allow small-scale communities to represent and communicate their views and daily lives on the web. Through the use of multimedia mobile phones, communities in different cities around the world have published images, videos and sound recordings in zexe.net for the last five years: taxi drivers in Mexico City, gypsies in Lleida and León (Spain), prostitutes in Madrid, handicapped people in Barcelona and Geneva, and motorcycle messengers (called motoboys) in Sao Paulo, Brazil. We call these web-based tools Community Memories, as they help communities represent and raise awareness about a commons. (Steels and Tisselli, 2008)

Collaborative tagging has become a crucial tool in zexe.net. Since the *canal*MOTOBOY project* `http://www.zexe.net/SAOPAULO` (2007 - ongoing), participants not only publish their daily experiences in the form of multimedia files, but they also tag them. Thus, zexe.net proposes a novel usage for tags by letting users assign them to what we could call "slices of life". In the following paragraphs, the basics of the zexe.net system will be described.

The zexe.net system comprises a set of PHP scripts and has a MySQL database as its backbone. The data structure used to store and organize the contents sent in by the participants is very straightforward:

- Multimedia elements, called attachments, are bundled together in small packages, called *messages*.
- Messages, in turn, are bundled together in larger containers, called *channels*. A channel can belong to a single participant or to various ones. Collective channels are created in order to aggregate specific shared topics.

The participants send their contents directly from the phone to the zexe.net web page by using the MMS (Multimedia Messaging System) service offered by the wireless digital networks. This service has the capability of sending a multimedia message directly to an email address. Thus, all the *messages* in the system are, originally, email messages containing multimedia attachments sent to zexe.net. As a consequence, all resources in the system are unique and they are associated to only one user. Tag assignments occur at the level of a message, meaning that all the multimedia files included in one message will share the same tag assignment. The tags assigned to a message can be entered as a comma-separated list directly on the phone, or by using an online editing application after the message has been received.

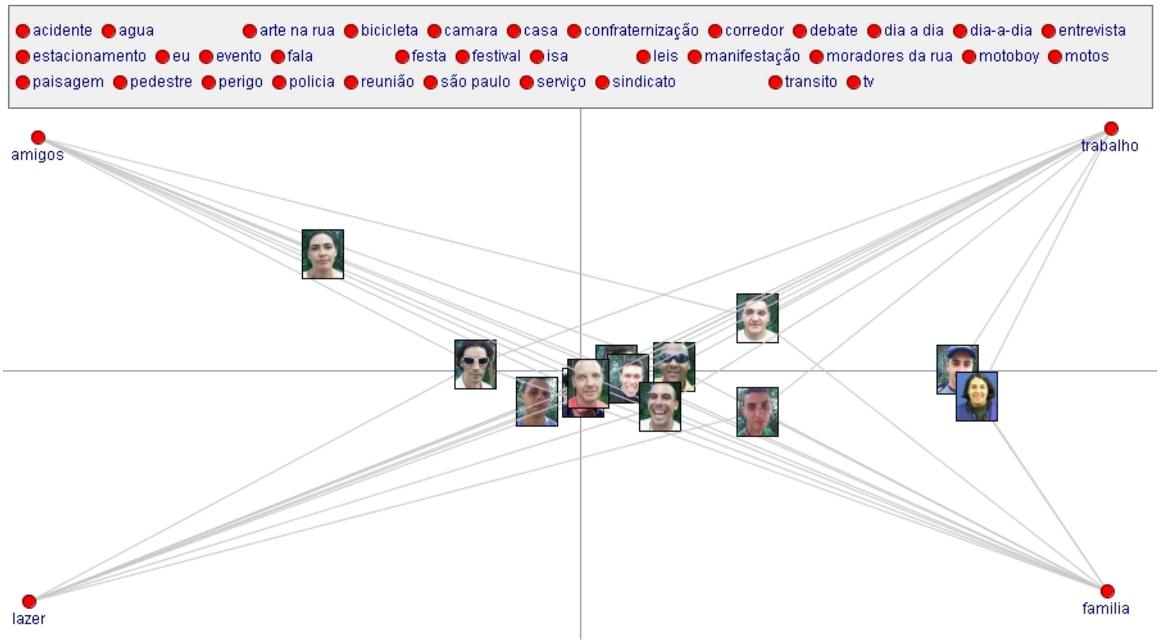

Figure 3.1: Tag-Participant Network in *canal*MOTOBOY*.

3.1 canal*MOTOBOY

In *canal*MOTOBOY* we have implemented an automatic detection of the singular and plural forms of the tags. These rules were written specifically for the Portuguese language. Thus, tags which exist in both of these forms are bundled together, publicly displaying only the most popular form. Each channel in *canal*MOTOBOY* features its own tag cloud. Tag clouds can be customized so that they emphasize either the frequency or the popularity of tags. Frequency refers to the number of times a tag has been used, either by a single user (in the case of individual channels) or by the whole group (in the collective channel). Popular tags can only be viewed in the tag cloud of the collective channel; their size is proportional to the number of participants who have used them.

Figure 3.1 shows the Tag-Participant network, a tool that shows the participants' position in relation to tags on a 2D plane. Tags are attractors, which means that the closer a participant is to a tag, the more he or she has used it. The objective of this tool is to reflect and compare the participants' tagging activities visually.

3.2 GENEVE*accessible

Tag cloud for frequent and popular tags can also be found in *GENEVE*accessible* `http://www.zexe.net/GENEVE`, the latest project at zexe.net (2008 - ongoing), which involves handicapped people who portray the state of the urban accessibility in Geneva, Switzerland. However, there are a number of significant innovations:

1. Every participant has a GPS-enabled mobile phone. Whenever the GPS is active, the corresponding geographical information (latitude and longitude) will be embedded in every photograph they take. Thus, this information is associated to individual multimedia elements. Google Maps is used for on-line GIS support.

2. Tag clustering was improved by introducing the possibility of manually creating groups of synonyms, which can include any number of tags. For example, the group "marches" can

Figure 3.2: Tag cloud showing highlighted (amigos, familia) and co-occurring (casa, ccsp) tags in *canal*MOTOBOY*.

contain its singular and plural forms, "marche" and "marches", and also typographic errors, such as "marhe". A second level of clustering is allowed by the possibility of creating channels from tags. Thus, for example, the channel "obstacles" includes the tag groups "dangers", "déviations", "entrées", "escaliers", etc.

3. Keyword searches for tags have been implemented.

4. A special application installed on the mobile phones provides the users of *GENEVE*accessible* with a list of the 10 most popular tags to choose from when tagging a photograph. This application is constantly updated with data coming from the zexe.net database. The application also allows users to add a tag which is not found on the list.

Because of the specific goals of *GENEVE*accessible*, an initial list of tags to be used for urban obstacles was convened with the participants at the start of the project. Needless to say, these tags quickly became the most popular ones.

3.3 The role of tagging in canal*MOTOBOY and GENEVE*accessible

In both projects, tag clouds act as search interfaces. The way in which tag clouds allow searches can be explained as follows:

1. A user select a desired tag *t1* from the tag cloud. The tag is highlighted, and the results are presented.

2. Only co-occurring tags are enabled in the tag cloud for further selection. If the user selects one of these tags, the search is refined by executing a database query that includes all the selected tags: *t1* AND *t2* AND *t3* AND ... AND *tn*.

3. The user can deselect any of the previously selected tags at any time.

4. To deselect all tags, the user can press the "reset" button at any time.

In *GENEVE*accessible*, tag searches not only produce a set of resulting photographs, but are also reflected in the corresponding map. The map itself becomes also an interface for navigation through the images, and can be used interactively, together with the tag cloud. In fact, selected tags or tag groups can be associated to markers of a specific color on the map.

Besides allowing publication and navigation, the zexe.net system also offers an online editing tool which participants use to manage their own channels and perform tasks on individual multimedia elements, such as editing their associated tags or their geographical location. The system also provides a control panel for the system administrator, in which different housekeeping tasks can be

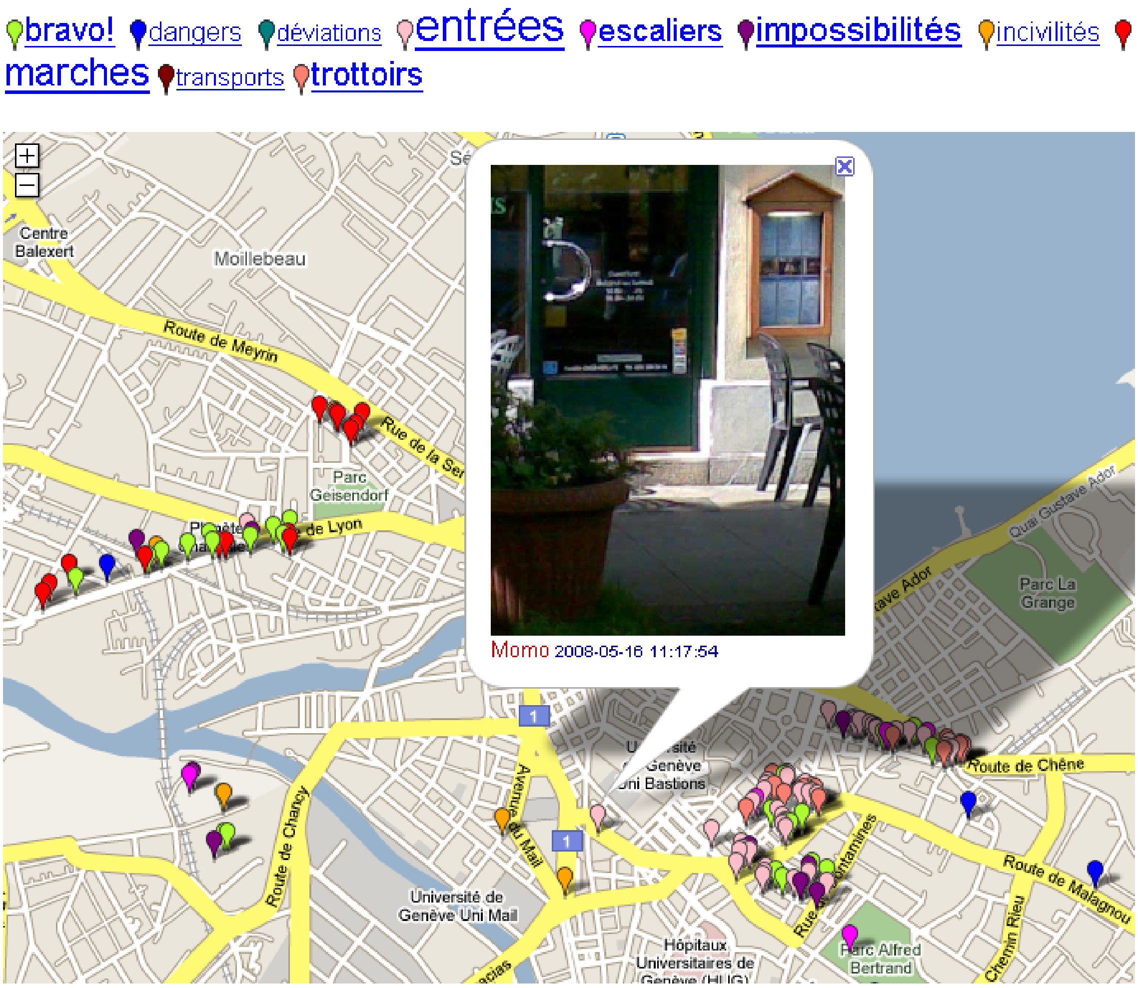

Figure 3.3: The "obstacles" channel in *GENEVE*accessible*, with the tag groups representing the different types of obstacles in the tag cloud. Each tag group has its corresponding marker color on the map.

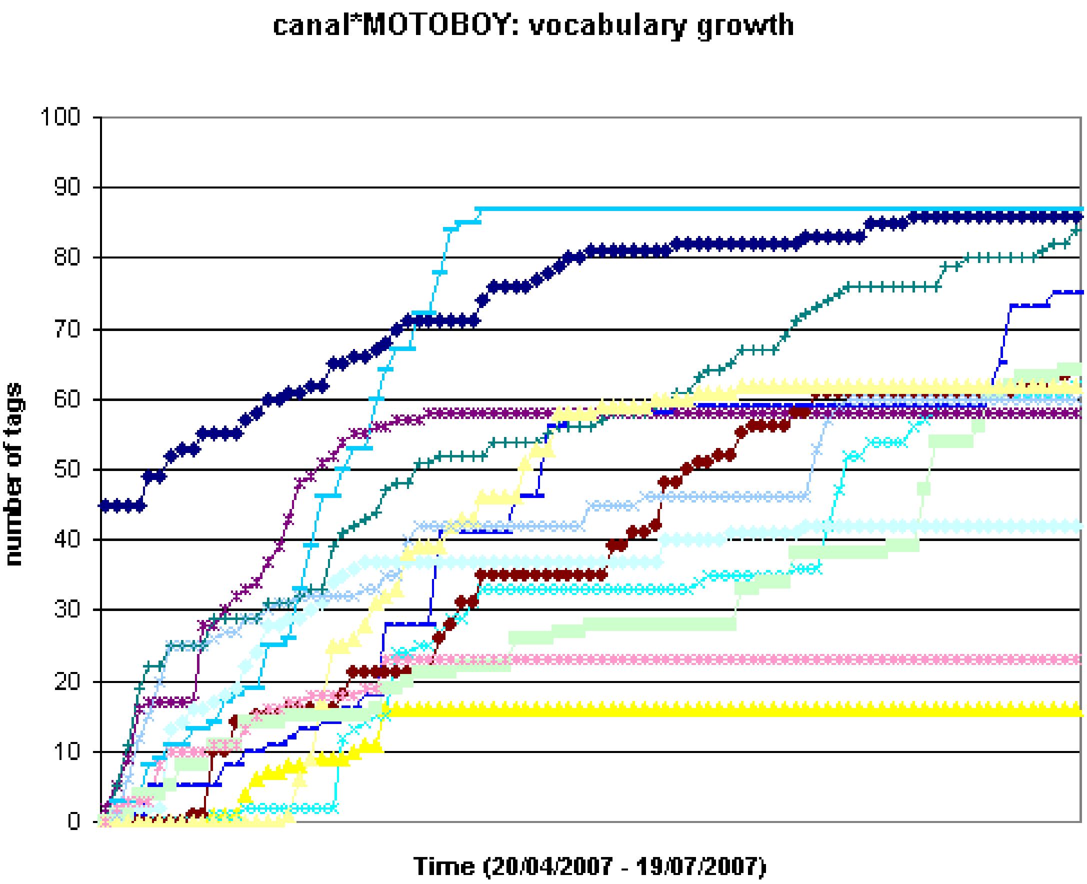

Figure 3.4: Vocabulary growth in the first three months of *canal*MOTOBOY*.

performed. Among these tasks are the creation of tag groups, the creation of tag-based channels and the choice of the specific tags to be highlighted on the map.

3.4 Analysis of tagging activity in canal*MOTOBOY and GENEVE*accessible

The following table represents the projects' tagging activity figures, as measured on 20/05/2008:

Project	Duration (months)	Users	Tags	Messages	Tag Assignments
canal*MOTOBOY	13	15	712	7.975	8.079
GENEVE*accessible	3	16	107	2.039	3.188

Vocabulary growth

These graphs show the vocabulary growth of both projects after a first three-month period. Each curve represents the growth of an individual vocabulary.

The following table shows the quantities of messages and tags existing in each project after these initial periods:

Project	Messages	Tags
canal*MOTOBOY	2.687	472
GENEVE*accessible	2.039	107

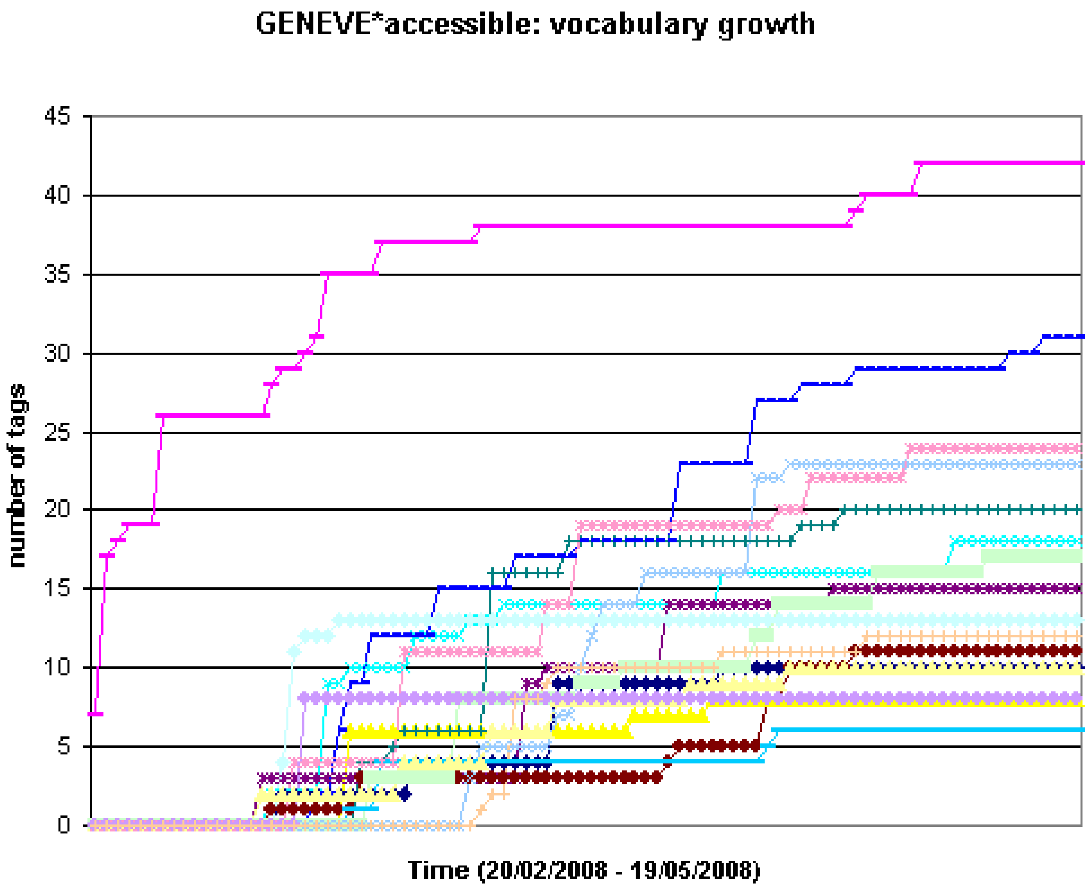

Figure 3.5: Vocabulary growth in the first three months of *GENEVE*accessible*.

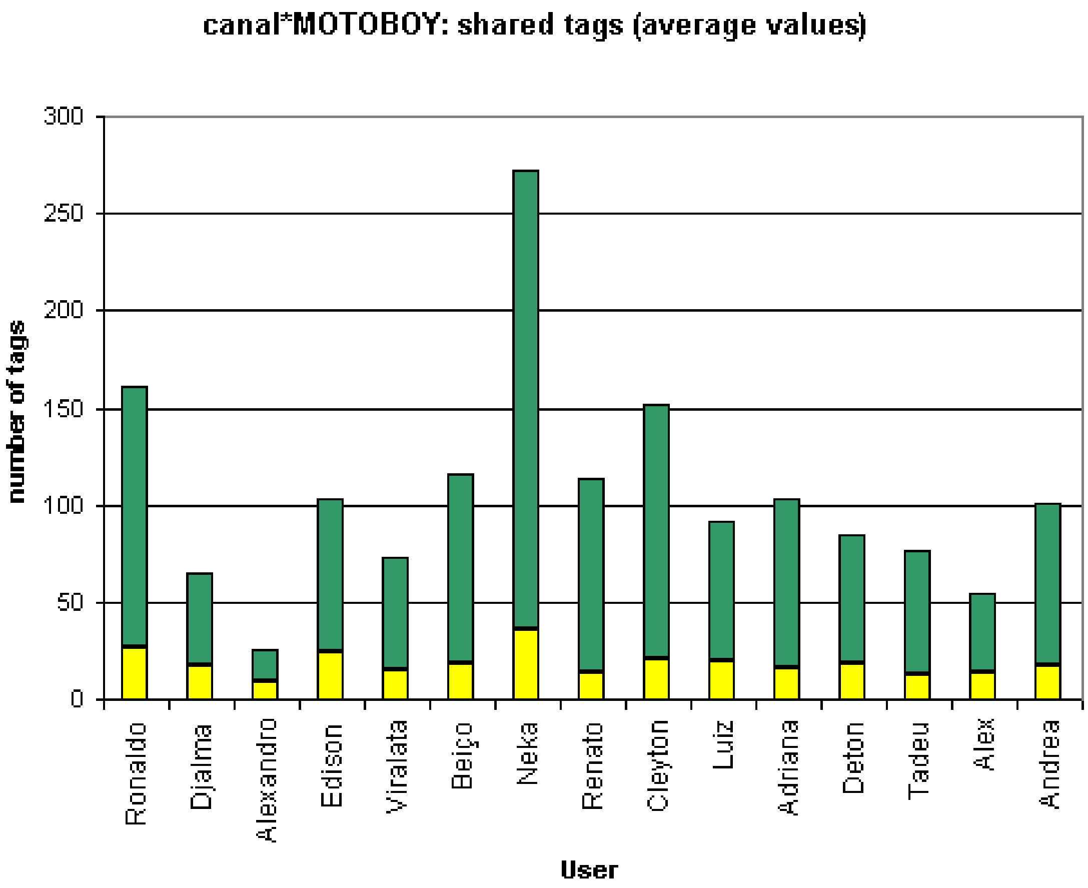

Figure 3.6: Number of distinct tags per participant (green) with average number of shared tags (yellow). Shared tags are tags which are also used by at least another participant.

It can be observed that, while the number of messages sent by these two groups are quite close, the number of tags in *canal*MOTOBOY* is more than 4 times higher than that of *GENEVE*accessible*. This can be due to the combination of two causes:

1. The *canal*MOTOBOY* project has open-ended goals, while *GENEVE*accessible* has very specific ones. The reduction in the scope of topics to be treated in *GENEVE*accessible* is reflected directly in the number of tags used by the participants.

2. The mobile application that lets users choose tags from a list is only available to the participants of *GENEVE*accessible*. Tag suggestion is very likely dampening the growth of individual vocabularies.

Alignment of tags in canal*MOTOBOY

These graphs represent the extent to which tags are shared in *canal*MOTOBOY*:

It must be noted that the motoboys form a very close group, and meet very regularly to discuss the project.

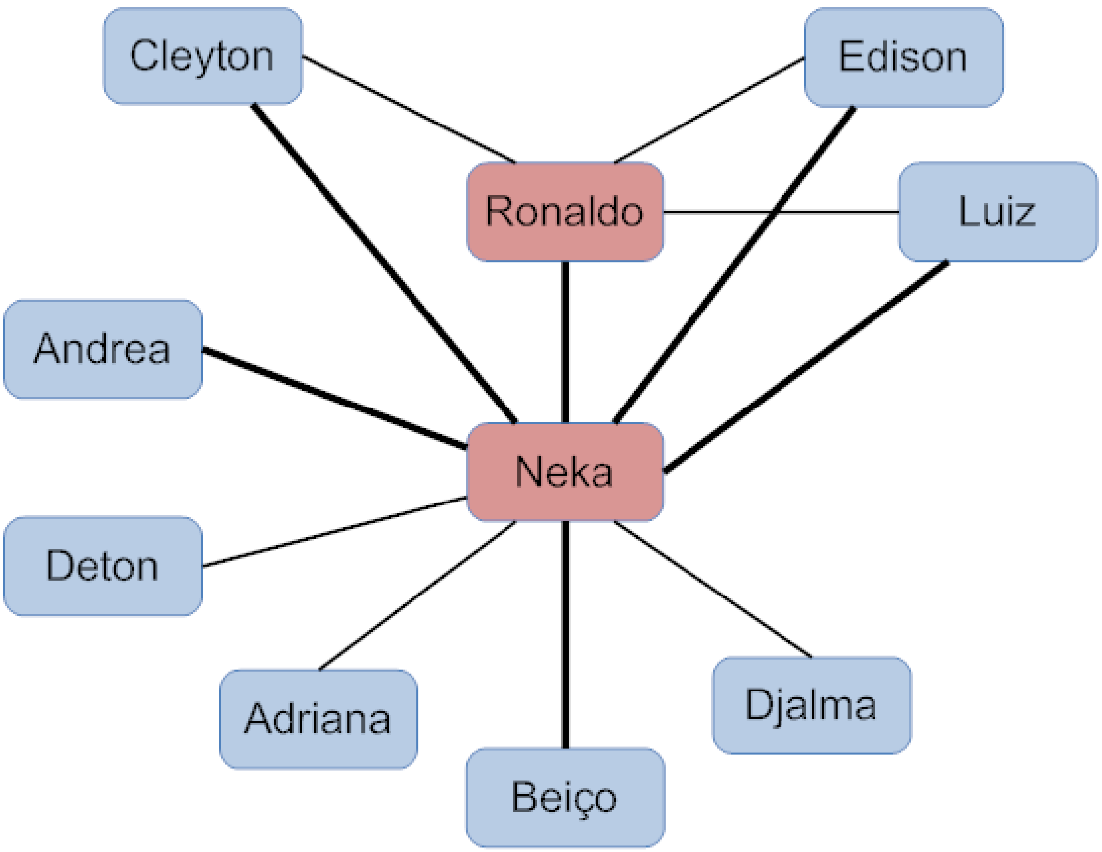

Figure 3.7: Network of participants who share more than 30 tags. Thicker lines represent more than 40 shared tags. Red nodes represent participants who have a degree > 4

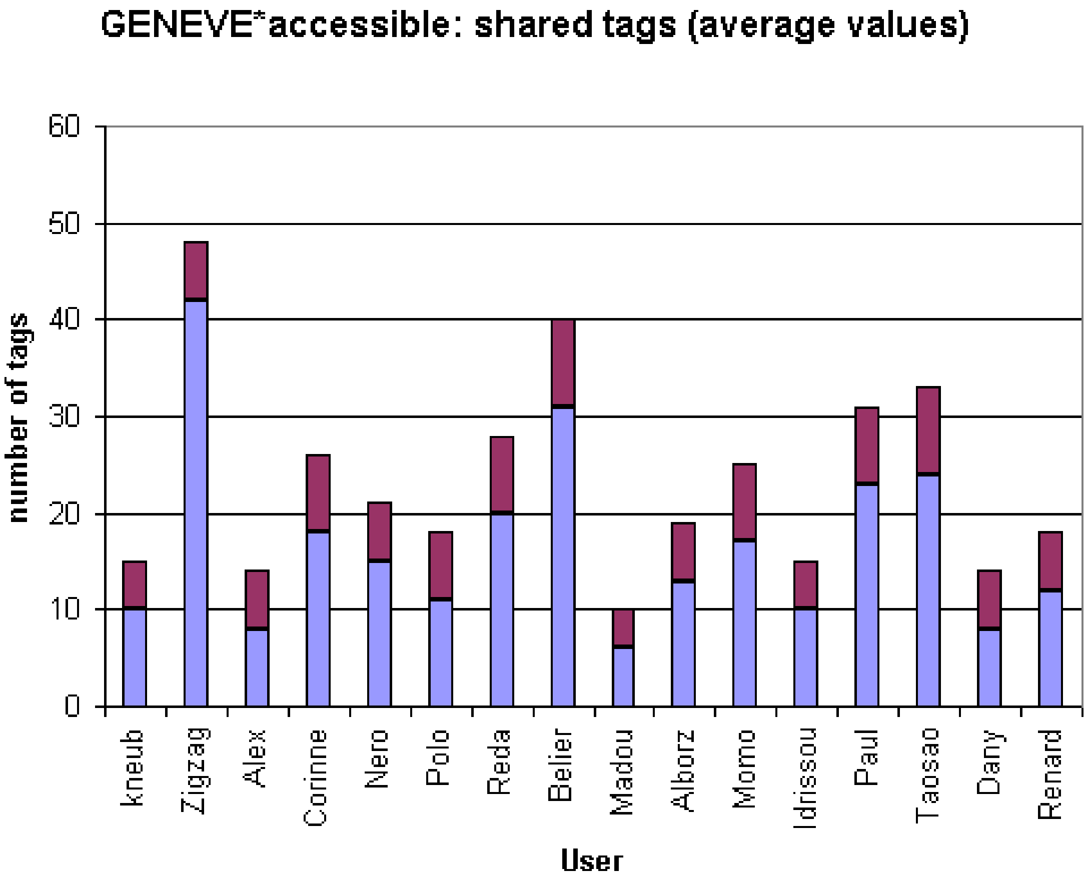

Figure 3.8: Number of distinct tags per participant (purple) with average number of shared tags (light blue). Shared tags are tags which are also used by at least another participant.

Alignment of tags in GENEVE*accessible

Individual vocabularies in the *GENEVE*accessible* project show a high degree of alignment, which can be due to the reduced scope of the project itself and the usage of the phone-based application which suggests a list of pre-existing tags.

3.4.1 Self-consistency in tagging behavior

The percentages of non-unique tags in the individual vocabularies of participants in both projects are illustrated in the following graphs:

The self-consistency measure represents the extent to which individuals are consistent with their own tags: a high percentage means that the participant regularly reuses her tags, while a lower one indicates that the participant normally uses unique tags. The causes which affect vocabulary growth in GENEVE*accessible also determine the comparatively high self-consistency rate of the individual users' tagging behavior in that project.

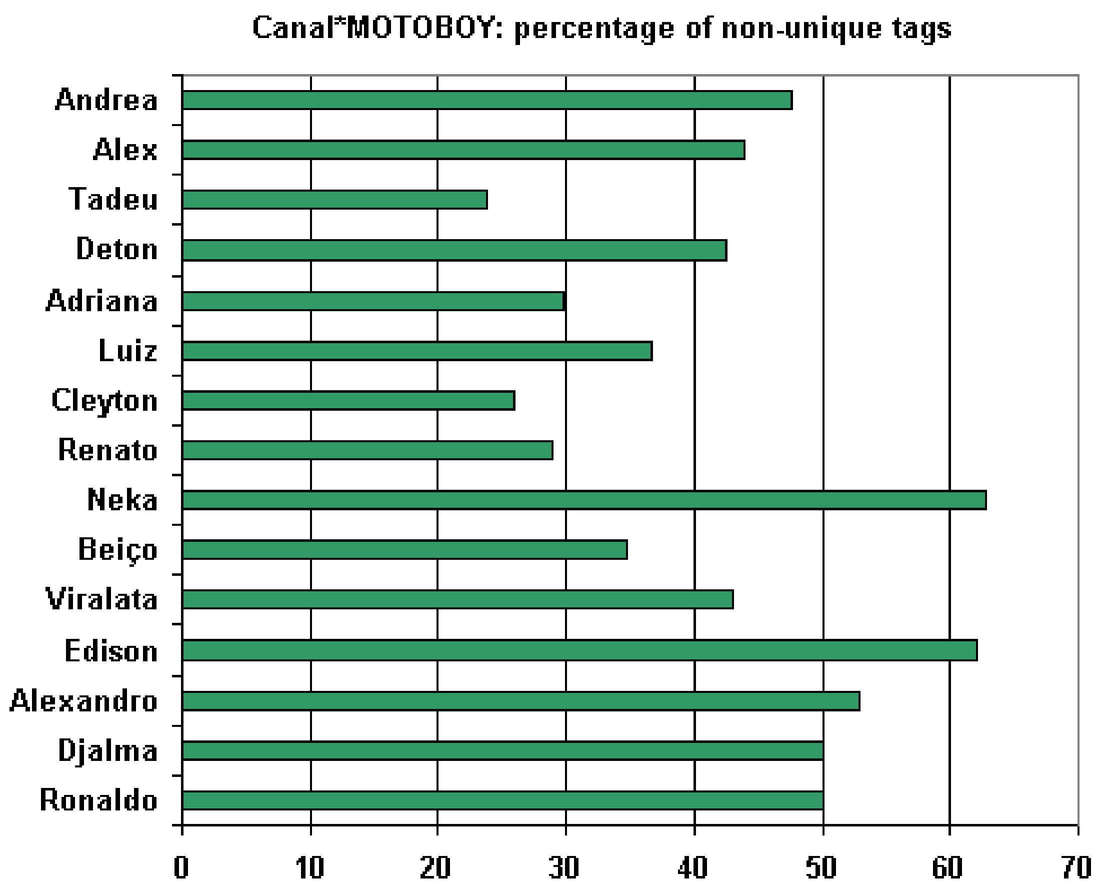

Figure 3.9: Self-consistency in tagging behavior, *canal*MOTOBOY*

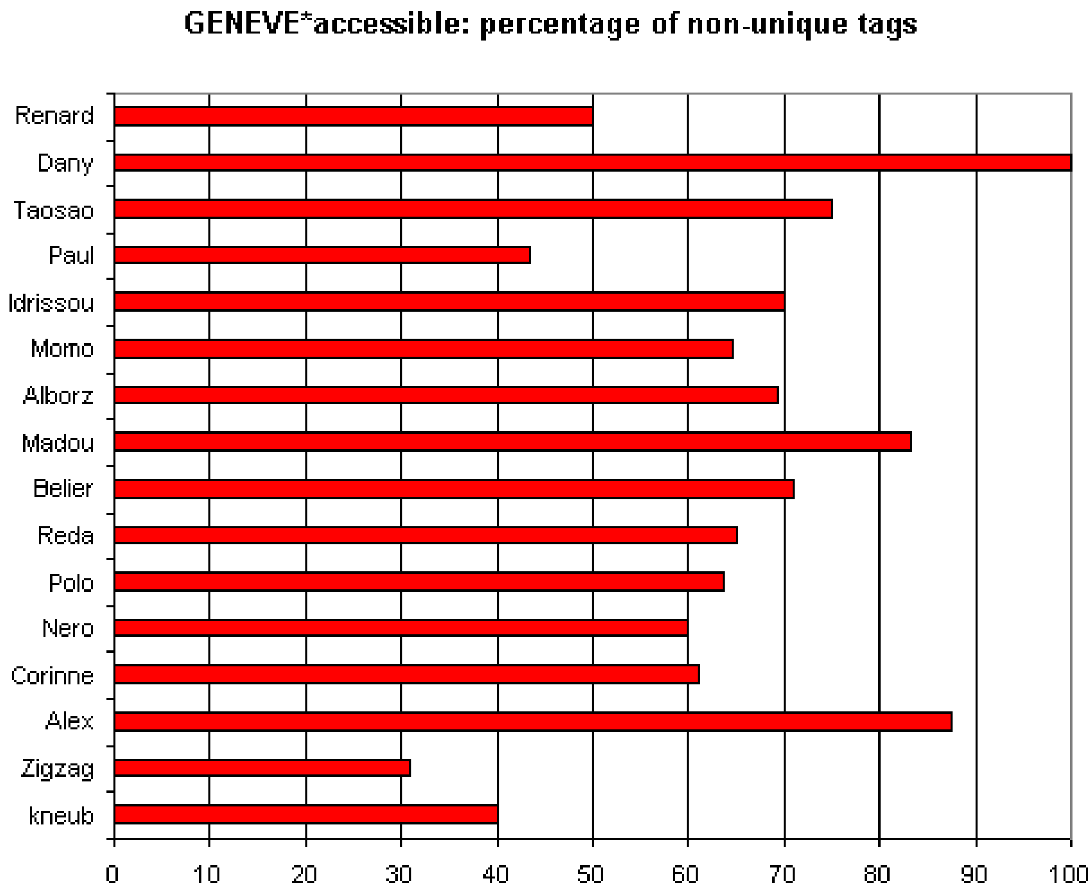

Figure 3.10: Self-consistency in tagging behavior, *GENEVE*accessible*

3.5 Conclusions

Even if the zexe.net system existed since 2003, the applications for the two latest projects mentioned here were totally re-written in order to support folksonomies. We found that the concept and the mechanics of tagging were understood very quickly by the participants of these projects, even by those who were not technically literate. The inclusion of folksonomies in these projects greatly improved the way in which the participants dealt with emerging topics, and provided a bottom-up way for representing the issues and views of the involved groups in a much more accurate and fine-grained way. By analyzing and comparing the tagging activity in *canal*MOTOBOY* and *GENEVE*accessible*, we also show how the scope of a project's focus and tag suggestion can influence the growth and diversity of folksonomies. The zexe.net system includes the basic functionalities of folksonomies: tagging with or without suggestions, tag clouds which are viewable using different criteria (frequency or popularity), filtering searches through tags and grouping. However, the system will evolve further in order to fit the necessities of each future project, and will incorporate the research results of the TAGora project as they are needed.

Chapter 4

Ikoru - A Test-bed for Collaborative Tagging and Content-Based Analysis

4.1 Motivation

The Ikoru system, developed at Sony CSL, is primarily used to experiment with collaborative tagging and content-based analysis. The project consists of a server-side component and a Web interface, which can be viewed at `http://www.ikoru.net`. Our motivations to develop this project included the following:

Content-based tools: Ikoru evolved from a research project that explored the combination of content-based analysis and collaborative tagging.

Data gathering: Running our own servers allows us to gather detailed user data and explore how the analysis of this data can improve tagging systems.

Extendible research platform: Ikoru aims to be an open platform that can be extended with new analysis and visualisation tools.

Multimedia: It was initially developed for images but has been extended to handle music files.

Small, reusable server: On the technological side, we have designed the Ikoru server as a small and stand-alone web component that can be easily reused and integrated in third-party projects.

We have made the first version of Ikoru available at the end of the first project year. In the second year we have extended the similarity search to audio. We have kept the Web site up and running since last year but we have been focusing more on targeted tagging experiments than on the growth of the Web site.

4.2 Similarity search for audio

In deliverable D3.3 (Section 3.3) we gave a very brief overview of a thorough study on automatic tag suggestion for music, based on the analysis of the audio signal. In that work, state-of-the-art classification techniques, developed at Sony CSL, were evaluated using a large music database that had been tagged consistently by a group of professionals using a well-defined taxonomy. Although the proposed classification technique perform better than existing techniques, the study also shows that the semantic inference remains extremely difficult. A direct translation of the used method from a clean-room database to an online tagging Web site is likely to yield unsatisfactory results.

However, content-analysis was the main reason for Ikoru's existence, and semantic inference one of it's aims. We showed last year that the image similarity search was a powerful tool, in particular when it is used with tags. Using tags, a visitor can narrow down the number of images that are displayed. At this point of the navigation in the archive, the simple content-based search becomes a useful tool to select a subset of the images. It can disambiguate, for example, between images of *apple fruit* and *Apple computers*, both tagged with *apple*.

This year we integrated the *contextual similarity search* for music into Ikoru and we decided to start with the simple approach that had been successful for images. The features we used are mostly those that are defined in the MPEG-7 audio standard[1]. To test it, we ran the feature extraction algorithms on the Last.fm data set, consisting of more than 18000 tagged music snippets of 26 seconds each. The similarity search can be tested on the Ikoru demo site (`http://demo.ikoru.net`).

4.3 Innovation through targeted experiments

At the outset of the project, we nourished the hope that Ikoru could grow into an active Web site. Despite the strengths of the system, this was somewhat wishful thinking. The reality is that in the last two years many sites have integrated tagging and that these sites can rely on considerable resources and infrastructure to continuously improve their offering. Technology transfers within Sony have been in principle possible and Ikoru has been presented to many product division within the group. However, the collaborations have been not trivial to set up because of the current tendency of Sony to outsource Web services.

As a results Sony CSL doesn't have a precise planning to promote Ikoru to a large audience. Instead, the current strategy is to continue to increase the usability and reliability of the software through its use in small but concrete projects. These focused projects can be managed much more easily and allow us to concentrate on innovative applications of tagging. In the future, we see Ikoru evolve as a generic back-end to store the information about resources, people, and tag assignments. We also see the focus of the tagging applications move away from purely Web-based applications towards real-world applications.

One such project is the artistic installation that is currently on display at the "You_ser" exhibition at the Zentrum für Kunst und Medien in Karlsruhe, Germany, and at the "Selective Knowledge" exhibition in Athens, Greece. Although this installation – a joined project with photographer Armin Linke – is a very particular use of Ikoru, it allowed us to gather a fair amount of data. More than 8000 visitors picked a selection of eight photos and tagged it using a special "editing table" designed for this purpose. The photos, printed in high-quality on solid boards, are taken from an archive of one thousand photos that are displayed on shelves in the exhibition space. Once the visitors tagged their selection, the editing table prints out a small booklet that they can take home.

Another interesting development, that has recently started is the use of Ikoru to store musical melodies (Pachet, 2008). Compared to photos or audio files, melodies can be analysed and generated at a semantically higher level. It has also the potential to reach a small but passionate community.

To facilitate such small tagging projects by other researchers and developers, and to let Ikoru evolve accordingly, we made the source code available under the GNU Library General Public License (LGPL). It can be found at `http://sourceforge.net/projects/ikoru/`.

[1]The feature extraction works as follows. The signal is split in overlapping chunks of 2048 samples (approximately. 46 msec long and 23 msec overlap). Each chunk is weighted by a Hanning window. For each chunk we apply a DSP operator such as, for example, the root mean square (RMS, related to energy level), the zero crossing rate (ZCR). We then calculate the first two statistical moments (mean and variance) to aggregate the values of each chunk into a single global value. Most operators return scalar values (RMS, ZCR, ...) except for Mel-frequency cepstrum (MFCC, 20-dimensional) and the Chroma analysis (measures the 12-tone distribution, 12-dimensional). The complete list of operators include: harmonic spectral centroid, harmonic spectral deviation, harmonic spectral spread, pitch, spectral centroid, spectral flatness, spectral spread, spectral kurtosis, spectral skewness, spectral roll-off, ZCR, RMS, RHF, HFC, IQR, centroid, harmonic spectral variation, MFCC, Chroma.

Chapter 5

TAGnet - a tool for awareness and management of personal metadata

5.1 Concept

People are not fully aware of the metadata they use to annotate resources. Tagging requires little effort and implies no strong committment to consistency: this fosters the externalization of large bodies of metadata, and at the same time makes the structure of the metadata rather unpredictable, even for the user performing the annotation. In the context of a single user, the tag co-occurrence network exposes many of the semantic relations among tags, and between tags and the broader context defined by user's interests and experiences. Visualizing the tag co-occurrence network and allowing the user to manipulate it provides her with a sort of "semantic mirror" that can be used for awareness, for navigation, and for re-organization of metadata. To this end, an application initially not envisioned in the work-plan, *TAGnet*, was designed to exploit the results of WP3 and WP4.

TAGnet is a prototype web-bases application (not yet available to the public) designed to provide users with a reflexive tool to expose regularities and patterns in their own tag-based annotations. Tagging patterns can reveal a lot about a user's experience, her interests and her emergent conceptualizations, but users are not aware of these patterns until these regularities are made explicit by means of data analysis and state-of-the-art visualization. TAGnet currently targets Flickr users and provides them with a view on their annotations (tags) that exposes actionable information.

5.2 Implementation

The development of TAGnet takes the move from an early experiment in this direction (`http://www.visualcomplexity.com/vc/project.cfm?id=231`), coded in Python and using Python bindings for the Flickr APIs (`http://flickrapi.sourceforge.net/`).

The very same logic and code were used to develop an interactive visualization by using state-of-the-state technology. In particular, the visualization client (Fig. 5.1) is a web-based application built on top of the Adobe Air framework (`http://www.adobe.com/products/air/`) and developed in Flex (`http://flex.org/`). The visualization and the force-based layout system use a customized version of the Flare (`http://flare.prefuse.org`) visualization toolkit. The Flickr APIs (`http://www.flickr.com/services/api/`) are used to fetch data from Flickr.

A Flickr user can connect to the *TAGnet* web site and insert her username. The system fetches from Flickr a list of the tags associated with each annotated photo and computes a co-occurrence network, which is subsequently visualized by using the above force-based layout engine. The user interface allows users to tune the number of tags displayed by the interface, and the threshold of co-occurrence controlling whether a link is drawn or not between two tags. The user can also

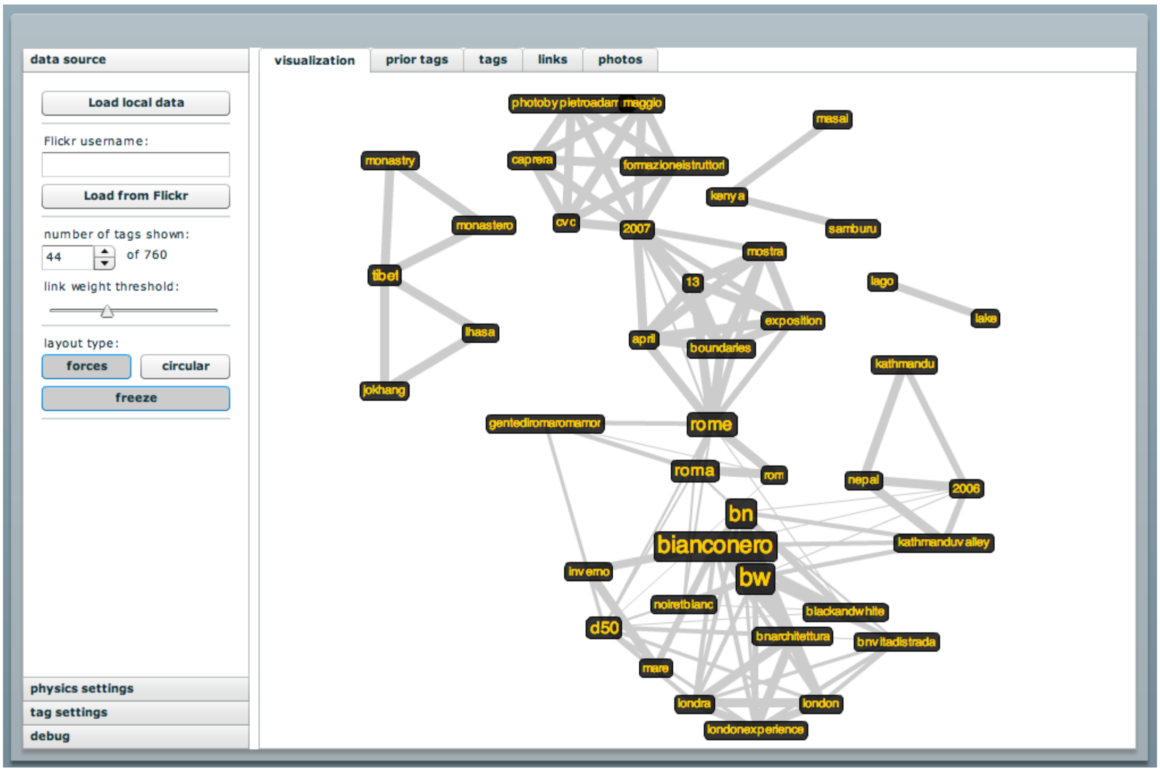

Figure 5.1: Search strategies in Bibsonomy

dynamically exclude tags, mark two tags as “synonyms”, mark a tag as a lexical variation of another tag, mark a few tags as “important”, and so on. By means of these actions, supported by the user interface, users can further structure the metadata and drive the visualization towards what they think is a better representation of the categories and conceptual structures they consider relevant. The user can switch among several different layout schemes and even freeze the layout engine to arrange tags manually according to her will. The resulting (user-manipulated) visualization of the tag co-occurrence network can be uploaded to Flickr and shared by clicking on a button.

5.3 Perspectives and road-map

In perspective *TAGnet* will be also used as a tool to explore emergent conceptualizations and tag ranking strategies for social annotations. To this end, extensive logging of interface events has been foreseen so that one can compare the measures of node importance and link strength computed by our system with the same notions as explicitly expressed by the user by means of the interface controls. This will yield insights into node ranking and similarity (Cattuto et al., 2008) in folksonomies, as well as a better understanding of what constitutes a better graphical layout (from the perspective of the end-user) in visualizing tag metadata.

In the long term, the user interface of the system will be cloned and customized to set up user studies targeting specific questions on user behavior, emergent categorization and conceptual structures, as exposed by the annotations of a given user. These experiments will be kept separated from the main system not to impair the applicative goal of *TAGnet*, which will be improved and kept focused as a tool for reflexive exploration of tagging patterns in the context of a single user.

Chapter 6

MyTag - Personalized Search and Exploration

6.1 Concept

Nowadays Web 2.0 platforms like YouTube, Flickr and del.icio.us provide large amounts of resources such as videos, photographs and social bookmarks. Common to the platforms is the classification by so called tags that can be used for organization and retrieval. A current limitation of tagging platforms is their confinement to a single media type. Furthermore, a magnitude of platforms exists for each media type. Thus, in both cases of searching resources of either the same or of different media type, a user has to search multiple platforms. For example, a user needs to search on del.icio.us, RawSugar and Bibsonomy to find bookmarks or on Flickr and YouTube to find media related to e.g. an artist. Another limitation results from the ranking of resources as implemented by platforms such as YouTube, Flickr, and del.icio.us. Usually, the overall popularity of a resource is used for ranking search results. A personalized search is currently missing that takes the interests of a user into account. MyTag aims at solving the previously described limitations of current tagging platforms by enabling cross-media search across images, video, and social bookmarks. It offers transparent access to different single-media platforms currently including Flickr, YouTube, and del.icio.us. Furthermore, it introduces personalization features such as a personalized ranking of search results. The architecture of MyTag ensures its extensibility towards further tagging platforms.

6.2 Implementation

MyTag is implemented using the Ruby on Rails framework as it supports efficient development of web-applications. The MyTag architecture realizes the model-view-controller paradigm (MVC). A view layer at the top is responsible for the interaction with the user while the control layer in the middle processes data from the model layer, e.g. by computing personalized rankings.

Two personalization features are provided for search: First, a search can be restricted to resources uploaded by the user. This feature requires that a user enters her external account names for Flickr, del.icio.us, and/or YouTube into her profile. Searching only own resources is implemented by using the corresponding feature from the integrated tagging platforms. The second personalization feature allows for ranking search results based on the user's personomy. The personomy is automatically built based on the resources the user picks from the result set. It is modeled by a vector p of tag frequencies representing the previous search interests of the user. As it is based on the implicit feedback given by selecting from the search results, no additional user effort is required to gain personalization. Using implicit user feedback is a very promising approach to personalizing search results or web browsing in general. This feature adds an advantage compared to systems

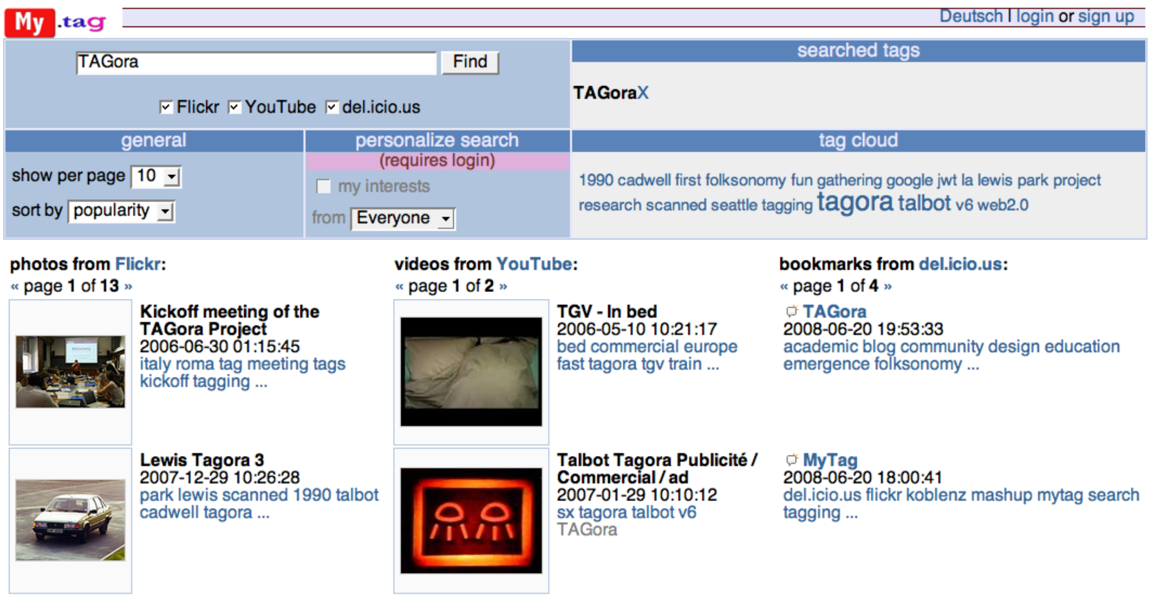

Figure 6.1: A screenshot from MyTag

such as Flickr and del.icio.us, where personalization requires adding resources to the system, i. e. the explicit feedback of users.

6.3 Perspectives and road-map

MyTag provides intuitive cross-media search and exploration over multiple tagging platforms. Personalized rankings of search results are featured by incorporating implicit user feedback gathered during the search activities of a user. As future work, we plan to evaluate and enhance the personalized ranking algorithm and compare it with state-ofthe-art approaches like FolkRank. Furthermore, we will focus on algorithms for merging search results from different platforms into a single result set.

Bibliography

Xwiki. 2004. http://www.xwiki.com/.

Ciro Cattuto, Dominik Benz, Andreas Hotho, and Gerd Stumme. Semantic analysis of tag similarity measures in collaborative tagging systems. In *Proceedings of the 3rd Workshop on Ontology Learning and Population (ECAI2008)*, 7 2008. URL http://arxiv.org/abs/0805.2045. http://arxiv.org/abs/0805.2045.

Martin Dougiamas. Moodle. 1999. http://www.moodle.de/.

Thomas Franz, Carsten Saathoff, Olaf Görlitz, Christoph Ringelstein, and Steffen Staab. Sea: A lightweight and extensible semantic exchange architecture. In *Proceedings of the 2nd Workshop on Innovations in Web Infrastructure. 15th International World Wide Web Conference (Edinburgh, Scotland)*, 2006. URL http://www.uni-koblenz.de/~saathoff/publications/sea.pdf.

Olaf Görlitz, Sergej Sizov, and Steffen Staab. PINTS: Peer-to-peer infrastructure for tagging systems. In *Proceedings of the Seventh International Workshop on Peer-to-Peer Systems, IPTPS08*, Tampa Bay, USA, February 2008.

Andreas Hotho, Robert Jäschke, Christoph Schmitz, and Gerd Stumme. Information retrieval in folksonomies: Search and ranking. In *Proceedings of the 3rd European Semantic Web Conference*, volume 4011 of *LNCS*, pages 411–426, Budva, Montenegro, June 2006. Springer. ISBN 3-540-34544-2. URL http://www.kde.cs.uni-kassel.de/hotho/pub/2006/seach2006hotho_eswc.pdf.

Robert Jäschke, Miranda Grahl, Andreas Hotho, Beate Krause, Christoph Schmitz, and Gerd Stumme. Organizing publications and bookmarks in bibsonomy. In Harith Alani, Natasha Noy, Gerd Stumme, Peter Mika, York Sure, and Denny Vrandecic, editors, *Workshop on Social and Collaborative Construction of Structured Knowledge (CKC 2007) at WWW 2007*, Banff, Canada, 2007. URL http://www2007.org/workshops/paper_25.pdf.

F. Pachet. Description-based design of monophonic melodies, 2008. Submitted to Computer Music Journal.

L. Steels and E. Tisselli. Social tagging in community memories. In *Proceedings of the 2008 AAAI Spring Symposium*, Social Information Processing, Stanford University, California, USA, 2008.

Zope. Kebas data. May 2002. http://sourceforge.net/projects/kebasdata/.