Functionality in a second generation tag cloud

Mogens Nielsen

Master's Thesis
Master of Science in Media Technology
30 ECTS
Department of Computer Science and Media Technology
Gjøvik University College, 2007

Avdeling for
informatikk og medieteknikk
Høgskolen i Gjøvik
Postboks 191
2802 Gjøvik

Faculty of Computer Science
and Media Technology
Gjøvik University College
Box 191
N-2802 Gjøvik
Norway

Abstract

With the growing popularity of collaborative tagging systems — like Flickr [1], del.icio.us [2], Technorati [3] and many more — the use of tags has become evermore popular. The need for a way to visualise the information within a collaborative tagging system has likewise become increasingly necessary.

The tag cloud is a popular way of visualizing the information in a collaborative tagging system. In this thesis, we have pointed to limitations of the first generation tag cloud as an information visualization interface. By examining the users of a tag cloud and the way they employ it to retrieve information, we have examined these limitations. And by studying related research, we have found ways to avoid these limitations. Some of these solutions has been implemented in a proof of concept to visualize a concept of a second generation tag cloud.

To evaluate the solutions and their impact on the tagging system and tag cloud-interface, we have examined their contributions in relation with a set of defined user-roles.

Our results show how a second generation tag cloud may look and how new functionality can improve it as an information visualization interface.

Sammendrag

Tagging er en form for metadata som har blitt mer og mer populær. Flere og flere nettsider bruker tagging for å beskrive ressursene som finnes i deres system. Disse ressursene kan være bilder på Flickr [1], eller andre nettsider som på del.icio.us [2]. Behovet for å visualisere informasjonen i et tagging-system har økt sammen med populariteten av disse systemene.

Tagskyen er en populær måte å visualisere denne informasjonen. I denne thesis har vi funnet begrensninger ved en første generasjons tagsky. Ved å undersøke brukerne av en tagsky og måten de bruker den til å finne informasjon, har vi undersøkt disse begrensningene. Og ved å studere relatert arbeid, har vi funnet måter vi kan unngå disse begrensningene. Noen av disse løsningen har blitt implementert i et proof of concept for å vise hvordan vi ser for oss en andre generasjon tagsky.

For å evaluere løsningene og deres innflytelse på taggesystemet og tagskyens interface, har vi undersøkt deres bidrag i forhold til et sett med definerte brukerroller.

Resultatene våre viser hvordan en andre generasjons tagsky kan se ut og hvordan nye funksjonaliteter kan forbedre den som et information visualization interface.

Preface

The author would like to thank Prof. Rune Hjelsvold for his support during this thesis.

Mogens Nielsen, 2007/06/30

Contents

1 Introduction

1.1 Tagging — Describing information

Metadata is information about data. Metadata describes a piece of information and thus makes it easier to retrieve from a larger collection of data. In the last years we have seen a growing interest for a type of metadata called tags. These tags are keywords or terms assigned to a piece of information by a either the author who created the information, the owner who shared it with a community or the end user who read it. The task of assigning these tags is called tagging and the user doing the tagging is referred to as a tagger.

When a tagger tags a piece of information — or a focus — she is describing this focus with a set of concepts that she relates to the focus. The focus can be any piece of information, on Flickr [1] the focus is a photo and on del.icio.us [2] (from now on referred to as Delicious) the focus is a web page.

The terms the tagger assigns to the focus summarizes what she thinks of when looking at e.g. a photo. By describing the photo the user allows herself and others to find the same photo with the help of a set of descriptive terms. We can describe tagging as filtering; from all the photos, show me the ones tagged with the word "jaguar" [4].

Figure 1 (from Rashmi Sinha's web log [5]) shows the process of tagging a focus. The user is observing a focus and as she is doing so one or more concepts related to this focus is activated. As opposed to hierarchical categorization where the user has to choose one of these concepts as the topmost category, a tagger merely has to assign the activated concepts to the focus. Wu et al [6] refers to tagging as social annotation and states that the reason to its popularity and ease of use is that:

> "Social annotations remove the high barrier to entry because web users can annotate web resources easily and freely without using or even knowing taxonomies or ontologies."

When several users have tagged the same focus with their individual concepts, they have created a semantic field or semantic landscape consisting of several terms and keywords that describe the focus in question. And as more foci are added we end up with a

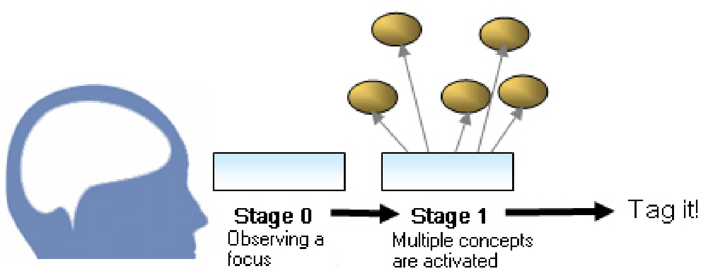

Figure 1: The process of tagging — from Sinha's web log [5]

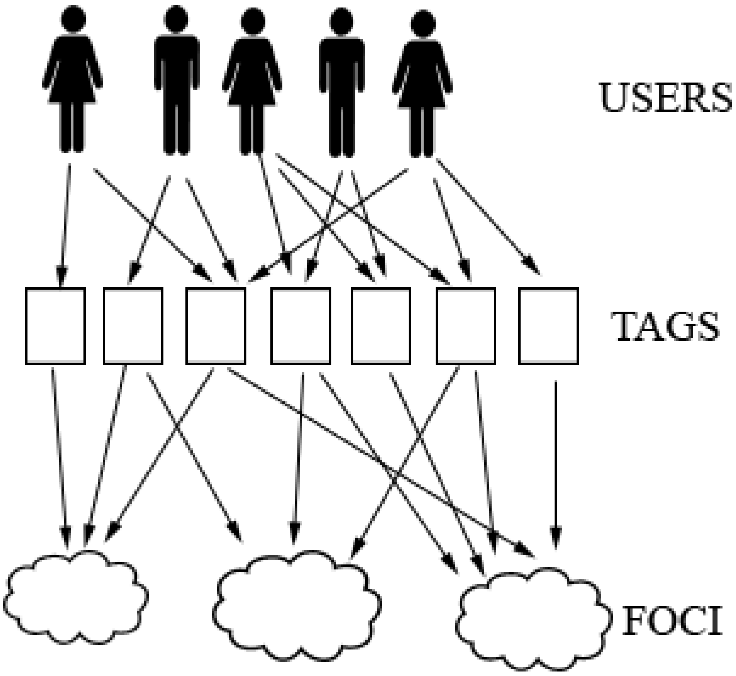

Figure 2: A model of a small semantic field

semantic field of foci and tags related to the foci. Each tag-focus-relation is created by a tagger and each tag can be assigned to different foci. To summarize, the semantic field contains a number of foci, tags and the relations between them. Some tag-focus-relations will occur several times as different users choose the same tag to assign to a focus.

In figure 2 we have visualized a semantic field consisting of a number of users, tags and foci. The figure visualizes how a user utilizes one or more tags to describe a focus. When more users they may tag the same focus with the same or other tags. And as more users add more foci to the system, the system grows in size of users, tags, foci and the relations between them all.

This semantic field can also be seen as a description of how different users perceive a focus. By choosing a specific focus and collecting all the tags assigned to that focus, we would end up with a list of tags and a frequency of how often each tag has been assigned to this focus. Thus a person that is unfamiliar with focus “y” could, with the help of the tags related to that focus, gather that focus “y” can be described by the tags “a”, “b” and “c”. And if focus “y” has been assigned the tags “a” and “b” several times, the user could infer that the tags “a” and “b” are good descriptors of the focus.

The problem is that when the semantic field around a focus is very large and consist of many different tags, it can be very time consuming and cumbersome to read through all the tags — it might just be easier to read through the web page. To help this new user gain a quick and easy understanding of what the semantic field comprises, the tag cloud has been put to use. The tag cloud will be more thoroughly explained in the next section.

The ease of tagging [5] is probably why it has become so popular on the Internet. Instead of having to choose a category — in the hierarchical sense — for a piece of information, the user can just tag that information with whatever words come to mind.

Figure 3: Jaguar - The cat

Furthermore the user does not have to consider the consequences of the tagging that one might have to when categorizing [5]. If a user was to categorize a photo of the South American cat, the Jaguar (figure 3), she would have to decide which category to choose for the photo and in which order. Should she categorize as a “cat” because it is a cat or as a “predator” because it is a predator? When tagging this photo the user would not have to consider these questions because tagging is not hierarchical. In tagging a focus can be members of more than one “category” — i.e the picture in figure 3 is both a “cat” and a “predator”.

As said by Rashmi Sinha in her web log [5]:

> “[. . .] the beauty of tagging is that it taps into an existing cognitive process without adding much cognitive cost. At the cognitive level, people already make local, conceptual observations. Tagging decouples these conceptual observations from concerns about the overall categorical scheme.”

1.2 Folksonomies — The language of tagging

On the bookmark sharing site Delicious, users who are tagging a new bookmark are presented with popular and recommended tags for this bookmark. In figure 4 we can see the web page that is presented to the user when she posts a bookmark to the system. She is presented with a list of tags that other users have used to describe the same website (popular tags) and in the list of recommended tags she is shown a combination of her own tags and those used by others. By looking at these tags new users are enabled to reuse the same tags when they are tagging a website instead of coming up with new tags — or reinventing the wheel as it were. This collaboration between users enables them to eventually agree on a common language in which to tag the resources. This is referred to as collaborative tagging and the common language is often referred to as a folksonomy [4].

By allowing end-users to see which tags other end-users have assigned to the same resource and by allowing the same end-users to change their tags the community in a collaborative tagging system will stabilize toward a common language or a folksonomy. This fact has been found by several researchers and will be more thoroughly explained later in the thesis.

The term folksonomy is attributed to Thomas Vander Wal in a discussion on an infor-

Figure 4: Posting a bookmark on Delicious

mation architecture mailing list and is a combination of “folk” and “taxonomy” [7] (there is some discussion as to whether this term is accurate [7]). This thesis will use the term folksonomy to refer to the common language that often evolves in collaborative tagging systems.

Later in this thesis we will show related work on folksonomies and advantages and disadvantages with folksonomies.

1.3 Visualizing information

Information visualization is used to visualize a large body of information so that it will be easier for an end-user to understand what that body of information is about. Such visualizations may be a subway map, a pie chart or any other visualization that communicates a large amount of information to a spectator in a way that makes this information easier to understand. By visualizing large amounts of data or information it is easier to obtain the information one is looking for and to present it to others.

In tagging systems the information to be presented is a large number of tags, consisting of words and terms, and the number of times each tag has been used. This information can be used to tell which tags are popular in a tagging community or it can help an end-user find information she is looking for.

When retrieving information it is common to differentiate between two methods of information retrieval (IR); finding and browsing. Finding refers to the method where an end-user looking for information is searching, often with the help of a specific query or a set of keywords. This query is entered into an information retrieval system and the end-user is presented with a result according to her query. Search engines such as Google [8] is an example of finding information.

Browsing refers to the method where the end-user is presented the available information on the system and she then looks through this information until she finds what she

Figure 5: A tag-cloud from Flickr

is looking for. When a person is walking through a library looking at books or looking through the file system on a computer, she is browsing.

Of course browsing can be a very cumbersome task when one is looking through a large collection of information and therefore information visualization is often implemented to help the user find what she is looking for. In accordance with browsing the semantic field of a tagging community one information visualization interface is much used, the tag cloud.

1.4 Tag clouds — What are they?

A tag cloud visualizes the semantic field by weighting the different tags according to how often they are used. This tag cloud is merely a weighted list of (most often) the most popular tags, often displayed in alphabetical order and their popularity visually weighted by font size [9]. This way the terms most often used to describe a focus is shown in a larger font size and thus making it easy for a new user to understand what she is looking at.

Figure 5 shows an example of such a tag cloud, this one is from the photo sharing site Flickr and shows the most used tags at the site. In this tag cloud some of the tags bubble to the surface of this tag soup — these tags are the popular ones and are therefore displayed in a larger font size. As one can see "party", "wedding", "Japan", "family" and "friends" are tags often used when users tag their photos.

Much like a table of contents in a book or a menu of categories on a web page, a tag cloud provide a means for users to form a general impression of the underlying set of content and a "gist" of what the web page is about [10]. The tag cloud helps unfamiliar visitors on a web page to form an impression about the contents of that site.

By considering the tag cloud a new user will quickly get an overview of which tags

Figure 6: A tag cloud of President Bush's 2007 State of the Union speach

are most used in this particular tag space and she can have a look at the different foci described by these tags. In the case of Flickr and Delicious, a user looking at a tag cloud can click on a tag in the cloud and she is then shown a number of foci (photos or web pages) which are tagged with that word.

Using a tag cloud to describe a tag space, as previously described, is one way to use a tag cloud. Another way to use tag clouds is to comprise a large text down to a simple tag cloud. TagCrowd [11] allows a person to create a tag cloud with the basis in a large text. By looking at that tag cloud a person who has not read the text can quickly make an estimate of what that text is about. As an example of this there was made a tag cloud of U.S. President George W. Bush's 2007 State of the Union speech [12] (figure 6). With the help of this cloud one can — without ever hearing the speech itself — estimate what the main context of the speech is because one can see which words where used most often (stop words such as "and", "of", "the" etc. have not been counted).

So why do we use tag clouds? As previously mentioned it is very time consuming and cumbersome to browse through a large semantic field and to make this process easier the tag cloud is used as an information visualization interface. The tag cloud is meant to give the information retriever (the end-user who is looking for a piece of information) a quick and easy overview of the information that is considered to be most useful to her. Most often this information is considered to be which tags are most popular at the moment.

1.5 Limitations of the first generation tag cloud

The tag cloud from Flickr (figure 5) and tag clouds like it are (in this thesis) referred to as first generation tag clouds. This is to differ them from what may come at the next stage; the second generation tag clouds.

With the help of such a first generation tag cloud, a user can quickly and intuitively get an overview of the most used tags in a tag space. A tag cloud is in a way like a satellite image of an area. With a satellite image of an area, one is provided with an overview of this area and for a closer, more detailed look, one can zoom in on the image. There are however some limitations to the first generation tag cloud and an overview of some of these are presented in this section.

We have divided the limitations in two groups. The first group is limitations that are inherit in the folksonomy and therefore are transferred to the tag cloud. The second group is limitations caused by the design of the first generation tag cloud.

1.5.1 Limitations inherited from the folksonomy

The limitations we have listed below are really limitations of the underlying folksonomy that the tag cloud visualizes. But because the tag cloud visualizes tags from a folksonomy, it inherits these limitations from the folksonomy.

No relationship between tags The alphabetical arrangement of the tags in the tag cloud does not infer any information about the relationships between the tags [9]. By arranging tags alphabetically inside a tag cloud, it is easy to find a previously know tag. But to a user who is unfamiliar with the tag cloud and is not looking for a specific tag, the lack of relations between the tags can be a problem. As the folksonomy is a flat structure, as opposed to a traditional taxonomy, this limitation is transferred to the visualization in the tag cloud.

Dominating topics In a tag cloud, one or more topics tend to dominate the cloud with their tags [13]. This is a problem because interesting tags are drowned out simply because the topic related to these tags is not popular enough for the tags to be visible in the cloud. The result is a tag cloud with many tags on the same topic, often synonyms. The problem of synonyms is inherited from the underlying folksonomies and will be more closely described later in this thesis.

1.5.2 Limitations caused by the design of the tag cloud

The limitations we have listed below, are caused by the design of the first generation tag cloud. The way the first generation tag cloud is designed, it produces some limitations, which we address as design-limitations.

No time control The first generation tag cloud shows us a view of the most popular tags in a tag space. The time frame for the presented tag cloud is predefined and often very long — e.g. the Flickr tag cloud in figure 5 visualizes the most popular tags over all time. In such a tag cloud a new tag will have great difficulty becoming visible because the popularity decay of the tag cloud is so short. Popularity decay is the rate in which something, in this case a tag, loses its popularity [14]. In the tag cloud from Flickr the tags have had such a long time to gather popularity that a new tag will not be shown — even though it is very popular at the moment. Such a static tag cloud does not give a view of evolution of popularity of tags in the tag space; we can not see which tags have increased or decreased in popularity the last days.

No detailed view In [15] Martin Graham mentions Ben Shneiderman's visual-information-seeking-mantra as "overview first, zoom and filter, then details on demand". The "overview first" part is covered by our first generation tag cloud. The "zoom and filter" part is however not. The most popular tags in a cloud is very visible to the user, but she does not have the possibility to get a more detailed view of the less popular tags.

Very wide search-field When a user finds a word in the tag cloud that she finds interesting, she clicks on this word. As she clicks on the word, the user is presented with all the resources that have been tagged with this word. As a search query, one word is a very wide search, and the result is a very large result set. In this result set, there will be many resources that are not related to the what the user was looking for.

No interactivity In the first generation tag cloud there are not many ways a user can manipulate the presentation of the tag cloud. At the moment interactivity in a first generation tag cloud is limited to the action where a user clicks on a tag and is presented with the related resources. There is no possibility for the user of a tag cloud to change the way the cloud is presented. The user viewing the tag cloud has to accept the organization of the tags in the cloud, the font sizes, colours and so forth.

2 Research questions and method

This chapter will describe the research questions and method for this thesis. The purpose of this thesis is to provide an overview of some of the limitations of the first generation tag cloud and show suggestions to solving these limitations. During the process of writing this thesis it became clear that it would be interesting to implement some of these suggestions in a proof of concept to show how a second generation tag cloud can be implemented. For the development of this proof of concept we chose two limitations from the first generation tag cloud and discussed these separately before they where implemented in our second generation tag cloud.

RQ1: How can the first generation tag cloud be altered to lessen the limitations that it inherits from the underlying folksonomy?

As we have previously stated, some of the limitations are really limitations of the folksonomy. This question will answer how the tag cloud can be altered to try and avoid these limitations. As well as how users may benefit from the alterations.

RQ2: How can time be utilized in a second generation tag cloud, and what are the consequences of doing that?

This question will look at why we should allow the tag cloud to utilize time and how this can be done. The consequences of this alteration of the tag cloud, will be discussed as to how they affect the tagging system and the graphical user interface of the tag cloud. The consequences will be evaluated to find out how they affect the simplicity of the tag cloud interface. And as to how the system design behind it can be designed to allow for time utilization.

RQ3: How can we enable the tag cloud user to get a more detailed view of the smaller tags, and what are the consequences of doing that?

RQ4: How can the tag cloud be altered to allow the user to narrow her search-field and manipulate the presentation?

When answering these two last questions, we will also evaluate their impact on the tagging system and the tag cloud interface. Specifically with regards to their impact on the simplicity of the tag cloud interface.

When answering research questions 2, 3 and 4, we will evaluate the proposed benefits of these functionalities, compared to how they may complicate the tag cloud interface.

2.1 Research method

The intention of this thesis is to gather information about the visualization of the semantic landscape in a tag space. By focusing on the tag cloud as an information visualization interface, this thesis will point to weaknesses in the first generation tag cloud. And with the help of existing literature, find suggestions on how to improve the first generation tag cloud.

To find limitations on the first generation tag cloud, we have studied related research. And we have read web logs on the Internet as a sort of “voice of the tag cloud users”. The web logs have all had either an opinion on the limitations of the first generation tag cloud or provided suggestions to future implementations. The literature has mainly been focused on folksonomies, tagging and tag clouds, in addition to literature on the development of information visualization interfaces. This literature has for the most part been acquired from the ACM Digital Library [16] and IEEE Xplore [17].

To evaluate the implementations of the suggestions found in this thesis, we will develop a proof of concept. The intention of this proof of concept is to analyse how second generation functionality can be implemented in a tag cloud. The proof of concept will also be used to evaluate the impact of the implementations as regarded to their functionality, and how they affect the tag cloud interface.

2.1.1 Developing visualizations

A tag cloud is an information visualization interface and this section will explore work related to the development of information visualization interfaces in the field of human-computer interaction (HCI). There is much research in both fields and this section will mention what is considered to be related to the further development of the tag cloud as such an interface.

Gershon et al states in [18] that

> ”Information visualization is about visualizing large amounts of information — it is the process of transforming data, information and knowledge into visual form making use of humans’ natural visual capabilities thereby enabling the user to observe the information.”

When developing interfaces for HCI it is important to have in mind that these interfaces will have to be useful for the people who are to use it. In [15] Graham et al discusses the development of visualizations for HCI. They state the problem that there is no universally agreed conceptualization or theory of HCI and hence no prescribed method for ensuring usability. However Shneiderman lists the eight golden rules of interface design in [19] (the rules are available online at [20]). As stated by Graham these rules are all very well, but knowing how and where to apply them, to develop better interfaces, is more problematic.

Through a system’s development and usage there will be discovered missing functionality or new demands to implementation — as is seen with the tag cloud today. Graham et al points to this problem with the task-artefact cycle — as seen in figure 7. The cycle illustrates how new artefacts (or systems) provide people with new possibilities that change the way people work which in turn determines their requirements of artefacts [15]. This reflects what is happening to the tag cloud; when its users are employing the tag cloud they discover new functionality which they want implemented. It is probable that the users of the tag cloud will be asking for more functionality at a later stage if its popularity continue like it is today. So how do we break out of this cycle? Graham’s suggestions lies in firstly taking a user-centred approach to design and secondly in recognizing that work will change as a result of our designs.

To try to combat this problem and gain a better “harmony” between the four facets of the overall work system (people, activities, technology and environments) Graham et al states in [15] that one can adopt a suitable development approach characterized as

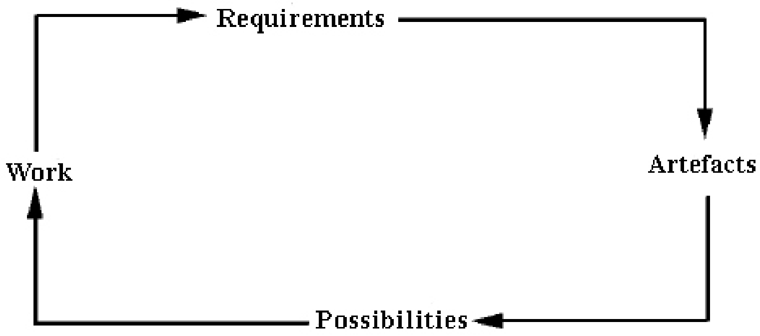

Figure 7: The task-artefact cycle from [15]

follows

- Take an *iterative* approach to systems development.
- Ensure real *user involvement* in the design process.
- Focus on the *work* to be done, not just the system functions.
- Attend to the *interface design* (including documentation).

Firstly Graham's list tells us to take an iterative approach to systems development. In this thesis we have defined the first iteration of the tag cloud as the first generation tag cloud and that the development of the tag cloud should enter its next iteration. In this iteration we need to evolve the tag cloud so that it may meet the demands of its users.

This brings us to Grahams next point; the involvement of the users into the design process. This thesis has found user interest into the development of tag cloud by researching web logs on the Internet. The people who use the tag cloud on a daily basis clearly have a meaning as to its usability and how it could be altered to better help them in their work. A second approach to ensuring user involvement is to perform user testing. As pointed out by both Rivadeneira et al in [10] and Halvey and Keane in [21] we have, despite the increasing popularity of tag clouds, seen very few studies that evaluate their effectiveness. This thesis has not performed any user testing, wee leave that for future work.

Graham's third point is that we need to focus on the work to be done and not just the system functions. This is exactly what this thesis intends to do namely bring into the open the need for a second generation tag cloud. With the help of this thesis it should be clear to systems developers, who want to implement a tag cloud in their system, that a tag cloud should not be just a static word soup. Instead it should be an information visualization interface that can visualize the information that the user wants to see. The user should not be bound to see which tags are most popular at the moment, arranged alphabetically. If the tag cloud is to be implemented in the solving of bigger problems than it is today, it needs to be able to visualize other information than just that. We have tried to focus on implementations that may improve the tag cloud by looking at the limitations of the first generation tag cloud and not just suggested new functions to implement.

The fourth and last point on Graham's list is that we need to attend to the interface design (including documentation). In the development of the proof of concept for this thesis we have also been thorough in the documentation of the system so that other developers can know how to develop a second generation tag cloud. And we have compared our proof of concept with that of a collaborative tagging system, so that we can find which design-differences are needed.

2.1.2 Evaluating tag clouds

A common evaluation measure for any technology is adoption by others [22]. Tag clouds have been adopted by many communities such as web loggers, photo sharing sites, bookmarking sites, news sites etc. The question that easily comes to mind is if tag clouds will stand the test of time or if they will wither away and be forgotten like so many new technologies.

Catherine Plaisant describes in [22] two information visualization system; Filmfinder and Spotfire. She goes on to describe how the development of these two system was an 11-year-long voyage and how the visualization technique based on dynamic scattergrams remained basically unchanged. The tool however has seen dramatic transformations. In the case of Spotfire there were designed separate products to fit the needs of specific application domains [22]. This thesis supports Kuo et al [23] and their beliefs for the future of the tag cloud:

> "Since the development of tag clouds is ongoing and their application to the summarization of text content is still not well understood, we expect rapid improvement and increased utilization in the future."

Plaisant [22] also describes TreeMap and DataMap and how they evolved from a prototype to an interface that is actually being used. The tag cloud is not a prototype anymore, but it needs to be more widely adopted and accepted so that more people can give their feedback and help with the evolution and improvement. In addition we should look at the potential to use tag clouds to solve bigger or different problems — the tag cloud is a tool for information visualization and information visualization is needed in accordance with other systems than just tagging systems. Only by investigating the tag cloud and its possibilities can we accomplish the perfect tag cloud. To quote Catherine Plaisant [22]:

> "We must complement our love for novelty by carefully integrating visualization tools into solutions for real problems [. . .] We should promote field studies, investigate new evaluation procedures and celebrate successes."

3 First generation tag cloud

This chapter will focus on the first generation tag cloud. We will look at the folksonomy that is underlying the tag cloud, and how the first generation tag cloud can be used for visual information seeking. We will examine the strengths and weaknesses of the folksonomy, and compare these to the use of the first generation tag cloud. By doing so we hope to clarify how the first generation tag cloud can be altered to avoid some of the limitations inherited from the folksonomy. With the help of earlier research on folksonomies, we will also show how the strengths of the folksonomy makes it a popular method for assigning metadata.

3.1 Tag cloud users

To understand the limitations and benefits of the first generation tag cloud we need to look at how it is being used and by what kind of users. For this purpose we have defined a set of user-roles according to what their intention with the tag cloud would be. The user-roles are defined in relation with the way they are looking for information and their familiarity with the system that employs the tag cloud.

Unfamiliar user An unfamiliar user is a user who is not familiar with the tag cloud or the underlying information of a specific system. This can be a new visitor to a web page like Flickr. The unfamiliar user is not looking for anything specific, but needs to form an impression of the underlying information.

Familiar user A familiar user is a user who is familiar with both the tag cloud and the underlying information. This user might be looking at the tag cloud to gain information about which topics interests the community in the system. The familiar user can for an example be a revisiting user to Flickr who wants to gain information about what has happened, in the underlying system, since her last visit.

Browsing user The browsing user is browsing through the information in a system like a person might walk through a library looking at books to find something interesting. The user might have a general idea of the information she is looking for, but has not got a specific term or word for it.

Searching user The searching user is retrieving information, but unlike the browser she has a specific term that she is looking for. Much like when a person searches on the Internet search engine Google [8].

Each of these users will have different needs in their usage of the tag cloud. And the two first user-roles (unfamiliar and familiar) can occur in conjunction with any of the two last, e.g. an unfamiliar, browsing user. Which gives us four user-roles comprised of the definition in the previous list.

- A browsing user who is *unfamiliar* with the system,
- A browsing user who is *familiar* with the system,

- A searching user who is *unfamiliar* with the system and
- A searching user who is *familiar* with the system

In our research questions, we stated that we want to evaluate how the users of the tag cloud can benefit from new functionality. This evaluation will be done with the help of these pre-defined user-roles. By looking at how these user-roles look for information in a tag cloud, we will evaluate how new functionality may help, or hinder, them in completing their tasks.

3.2 Retrieving information with tag clouds

This section will look at how the tag cloud can be used as an information retrieval interface, we will do this with the help of the previously defined user-roles.

Unfamiliar user In [10] Rivadeneira states that one of the tasks a tag cloud can support, is impression formation. A user looking at a tag cloud gains insight into which topics are popular, and to some extent, which topics are not that popular. In [23] Kuo et al found that a tag cloud is advantageous to the standard list in presenting descriptive information, and in reducing user frustration. In their study they found that the quality of the answers to descriptive questions (e.g. "is TFL a transcription factor?"), was generally higher with the use of a tag cloud as opposed to a list. As a tool for getting an overview the tag cloud seems to be a good implementation, and can therefore be useful to an unfamiliar user who has not visited the website before.

Familiar user A familiar user will already know what the underlying system contains, and is in need of more information than just to format an impression. The information the revisiting user can gather from the tag cloud depends on the time frame set for the tag cloud. If the time frame is very long the tag cloud will not have changed much and therefore does not provide the familiar user with any new information. The problem with the static nature of the first generation tag cloud has already been described and will be examined more closely later in this thesis.

Browsing user According to [24] browsing is one of the strategies most often cited as benefiting from visual support. According to Lin [25] browsing is a superior strategy when

1. there is a good underlying structure so that items close to one another can be inferred to be similar,
2. users are unfamiliar with the contents of a collection,
3. users have limited understanding of how a system is organized and prefer the less cognitively loaded method of exploration,
4. users have difficulty verbalizing their underlying information needs and
5. information is easier to recognize than to describe.

By comparing this list to the first generation tag cloud we can find drawbacks that should be limited to improve the first generation tag cloud as a browsing interface.

Structure Lin's first point is related to the lack of structure that we have already described. A suggestion to how we can solve this problem will be examined more closely later in this thesis. But Lin's point indicates to us that the lack of structure in the tag cloud inhibits it as a browsing tool.

Unfamiliarity As described above, the tag cloud is a useful tool for an unfamiliar user. And this point in Lin's list shows us that the tag cloud can be a very useful tool for an unfamiliar user. When visiting Flickr, a user knows that the contents of the site are photos, but that is it. And a new user to Delicious, does not know exactly what to find inside such a large collection of bookmarked web pages. As the user is about to embark on an exploration of this collection, it is very useful to for her to have a browsing interface to support her. We state that the first generation tag cloud is a good interface for this purpose.

Limited understanding One of the weaknesses of a tagging system is its lack of structure. In a taxonomy there is a strictly defined hierarchical structure where resources are located. In a tagging system there is no hierarchy; each resource is equally a member of the "category" represented by each tag. Lin's third point tells us that the tag cloud can be a useful tool when a user is browsing through a system where she is not familiar with the underlying structure. Even though neither the first generation tag cloud or the underlying folksonomy has a strict structure, there is a structure that the user is not familiar with. Later in this thesis we will show how a hierarchical structure, or clusters, can be inferred from a folksonomy.

Difficulty verbalizing Sometimes when a user is looking for information, she may have difficulty verbalizing what kind of information she is looking for — but she has a general understanding of what she is looking for. The user may also find it easier to recognize the answers she is looking for than to define her question. In his two last points, Lin points to these situations as situations where the user can benefit from visual support. We can compare this, again, to the person browsing the library; the person knows she is looking for information about architecture, but has difficulty defining a question that needs to be answered. Much like this person would be browsing the architecture department of the library, a tag cloud user can look at the words in the cloud to find a word or a term that seems familiar. This word can then be followed and the user may find what she is looking for.

All in all, our investigation of Lin's points seem to indicate that the first generation tag cloud is a good browsing tool. But it could perhaps benefit from some alterations.

Searching user A user that is looking for information by finding and not browsing, has a specific question that she needs answered. This question can often be answered if it is reduced to a query and entered into a search engine. In a tag cloud, it is not possible to search for a specific query. When a user wants information from a tag cloud, she finds an interesting word and clicks on it. She is then presented with resources that have been assigned this tag. This one-word-search-query is very rough and will return a lot of resources, many of which are not related to what the user is looking for. If a user is to use the tag cloud to *find* information, it needs to be more precise.

To summarize, the tag cloud seems to be a useful tool for a browsing user, in particular a browsing user who is unfamiliar with the system behind the tag cloud. However, in folksonomies, and therefore tag clouds, there is a lack of structure. The lack of structure is a limitation of the tag cloud that we have outlined previously in this thesis and a suggestion to combating this problem will be more closely described later in this thesis.

3.3 Folksonomies

The term folksonomies refer to the language that emerges in a collaborative tagging system (CTS). This language evolves as the different users of the CTS are adding, using and tagging the resources in that system. Because the tag cloud visualizes an underlying folksonomy, many of the limitations of the first generation tag cloud are in fact inherited from the folksonomy. In this section, we will look at research in the field of folksonomies to help understand the limitations of the first generation tag cloud. We will also look at the strengths of the folksonomy, to try and understand why it has become so popular.

There is quite a bit of research on folksonomies and Golder and Huberman show that on Delicious stable patterns emerge [4]. This fact tells us that in such a collaborating community, a folksonomy actually does develop. Golder and Huberman hypothesize that the users stabilize on a set of terms mostly because they are influenced by the tagging behaviour of the other community members. Halpin et al confirms this evolution towards a common language in [26], where they show that tagging distributions tend to stabilize into power law [27] distributions.

Metadata can be created in three ways; by professionals, by the content creator or by the end-users who are using the data on a daily basis [7]. Collaborative tagging is about letting these end-users create the metadata in collaboration with one another.

In the traditional method for describing data, one or more experts creates a thesaurus of categories and takes care of the indexing of the documents themselves [28]. This method is tedious and often requires serious training and education [7]. The benefits are that the insertion of erroneous data is limited, however the insertion of different viewpoints and aspects of the foci are also limited [28]. It is important to capture the different aspects of a focus when tagging, because different users have different educational and cultural backgrounds. To some people a piece of raw fish is a “futomaki” and to others it is just a piece of “fish” or “sushi” [13]. This issue has also been addressed by Golder and Huberman in [4] and is often referred to as the vocabulary problem [29].

Furnas et al showed in [29] that when two people are spontaneously assigning a term to an object, the probability for them to use the same term is less than 0.20 — i.e. if one person assigns the name of an item, other untutored people will fail to access it on 80 to 90 percent of their attempts [29]. This vocabulary problem is a problem with metadata created by a professional or the creator of the resource. When only one person chooses the terms to assign to a resource, the probability for another person to use that same word when searching is too small to make such a system user friendly. A practical example of how difficult it is for two people to agree on simple descriptive words for a photo, is the ESP Game [30] by Luis Von Ahn and Laura Dabbish [31]. To limit this vocabulary problem more people can be involved in the process of metadata creation. As more people are involved in the process, more facets of the tagged resource is included. However, the amount of erroneous tags used to describe resources in the system is likely to increase. Quintarelli argues in [32] that the trade-off between simplicity and precision makes sense in most practical cases.

The insertion of erroneous metadata is not the only contributor to the lack of precision in a collaborative tagging system. Sometimes the lack of properly describing metadata can add to this. Sen et al [33] discovered in their system MovieLens that a large part of their users did not tag their resources simply because they could not think of any tags to use. Offering tag suggestions such as Delicious does (see figure 4), may be a way to avoid this problem by helping the users. Sen et al also suggests to seed new tagging systems with a large set of tags of the preferred type and thereby hoping to limit the amount of unwanted tags in the system.

Adam Mathes [7] and Golder and Huberman [4] have pointed out other limitations to collaborative tagging systems. These are outlined below.

Ambiguity/Polysemy Ambiguity or polysemy describes the fact that one word can have several meanings. Adam Mathes [7] points to an example of searching for the tag "filtering" on Delicious and the results contained both Internet Radio, Bayesian Filtering and a guide to filtering vodka. Clearly the word "filtering" has different meanings to different people. This fact will generate more noise in the result returned to a user who is searching for information, as she is shown resources that do not relate to what she had in mind when querying the system.

Synonyms Synonyms are the opposite of polysemes and are multiple words that mean the same. Examples of synonyms are "student" and "pupil" as well as "interesting" and "fascinating". In tagging systems there is no synonym control (that the author is aware of) and this leads to users tagging foci with many different words which in fact mean the same.

Plurals and singulars are another example of synonyms. Users tend to use both the singular and plural term of a word when assigning it as a tag to a focus. As mentioned above, to limit this problem, Delicious offers related tags based on the tags that the tagger has already entered. The problem of synonyms in folksonomies has also been addressed by Li et al in [34].

Spaces and multiple words Both Delicious and Flickr seem designed primarily to deal with single words and not phrases [7]. Seeing as how it is a computer that interprets the tags and to a computer the term "collective intelligence" is not a phrase or a term, but two different words, terms and phrases will not be recognized. Some taggers avoid this problem by spacing words with the underscore "_" or just by dropping the space and putting the words together such as "opensource", "cameraphone" or "sometaithurts" ("so meta it hurts" on Flickr).

Previously in this section the limitation of the vocabulary problem was mentioned as one the strengths of a folksonomy. Adam Mathes points to more strengths of the folksonomy and these are outlined below.

Browsing and serendipity Browsing and its relation to folksonomies and tag clouds has already been described in this thesis. Adam Mathes states that there is a fundamental difference in the activities of browsing for interesting content, as opposed to direct searching in a query:

> "It is similar to the difference between exploring a problem space to formulate questions, as opposed to actually looking for answers to specifically formulated questions."

By browsing, users are also more likely to discover resources by serendipity. Lawrence Block [35] described serendipity as follows

> "Look for something, find something else, and realize that what you've found is more suited to your needs than what you thought you were looking for."

Adam Mathes describes how he himself discovered many resources, when researching for his paper, that he likely never would have been exposed to. This leads to new discoveries that might or might not give the researcher new input.

Serendipity has also been described as one of the strengths of folksonomies by Quintarelli in [32] and Montero in [9]. Li et al [34] states that folksonomy is not only a method for organizing content to facilitate the users who create it, but also a navigation mechanism for users to discover interesting resources.

Desire Lines In architecture, landscape designers sometimes refer to something called desire lines. Desire lines are the foot-worn paths that sometimes appear in a landscape over time [7]. They can also be described as the shortest route between an origin and a destination [36]. Many streets in old cities began as such foot-worn paths and evolved into streets. Peter Merholz notes in [37] that:

> "A smart landscape designer will let wanderers create paths through use, and then pave the emerging walkways, ensuring optimal utility."

One of the greatest strengths of a folksonomy is that it directly reflects the vocabulary of its users [7]. Much like desire lines in architecture shows how people travel from point A to point B, the folksonomy shows how the users of a collaborative tagging system wants to describe different resources. The emerging folksonomy in a collaborative tagging system could, if one wanted, be used as a framework to develop a controlled vocabulary — i.e. a vocabulary that speaks the users' language [7].

3.4 Summary

As we can see, the folksonomy has both strengths and weaknesses. Many of the weaknesses in the folksonomy are transferred to the tag cloud, and thereby limiting it as an information visualization interface. Problems such as synonyms are passed to the tag cloud, leading to the domination of a few topics; one topic will have many related words, who nearly "takes over" the tag cloud. The effect of these limitations could perhaps be lessened by changing the process of tagging a resource. This thesis will not be looking at the folksonomy or the process of tagging resources. We will focus on the tag cloud and how new functionality can limit the problems from the underlying folksonomy.

The strengths of the folksonomy, makes the folksonomy a possibly very lucrative method for assigning metadata. Some of these strengths however, are not very well taken care of in the first generation tag cloud. In this chapter, we have described how browsing for information can benefit from a folksonomy.

By relating the information from this chapter to our previously defined user-roles, we will try assess how each of them can benefit or be hindered by the folksonomy. Because we are not looking at the first generation tag cloud, but at the folksonomy, we will not be considering the user's familiarity to the system, but merely whether she is looking for information by browsing or finding.

The browsing user will be "walking" through the system looking at topics, and the topics are listed as keywords or terms. This process is much like walking through a library and reading the titles of the books. While the searching user has more of a definition of what she is looking for, and where to find it, the browser is just looking for topics that seem interesting to her. The flat and unstructured nature of the folksonomy may become frustrating to the searching user, because she can not find the parent- and child-categories that she is looking for. Because of this lack of structure, looking for resources on a specific topic in a folksonomy, can be much like trying to find books on programming in a library where all the books are placed randomly amongst the shelves. This indicates that a folksonomy may be more suitable for browsing than for finding.

As we have previously explained, the folksonomy catches many facets of each resource. Each facet has been assigned by a different user, and that user may have conceived the resource differently than you or I have. So when a user finds a term that seems interesting to her, that term may be tagged to resources that she would not have found in a hierarchy. In a hierarchy, a leaf-node can be located by travelling from the root-node and through the branches and down to the leaf-node. This way, the distance between leaf-nodes can be very large even though the resources in them are related. In a folksonomy there is no hierarchy, and each resource is a member, of the "category" represented by each of the assigned tags. This lack of structure in the folksonomy, may hinder a user who is looking for information by finding. But it may also help her find new resources by serendipity. The same can be applied to the browsing user.

The problem of polysemes and ambiguity was mentioned earlier. This problem is applicable to all users of the folksonomy. When looking at the tag "filtering", the user has no way of knowing what kind of filtering this word refers to. With a structure in the folksonomy, it may be easier to understand what these polysemes refer to. With cluster, the user can look at which words are related to the "filtering" tag. And with a hierarchy the user can look at the "parent"-tags of the tag she is looking at to gain more information.

In this chapter we have been looking at the first generation tag cloud in relation to our four user-roles. Our research indicates that the tag cloud may not be suitable for a user who is looking for information via finding. Instead, it indicates that the tag cloud is a suitable tool for finding information via browsing. The tag cloud does have some restrictions that, if limited can improve it as a browsing tool.

Our research also indicates that the tag cloud is a very suitable tool for a user who is unfamiliar with the underlying system. For an unfamiliar user, the tag cloud helps her form an impression of what she can find inside this system. The same goes for a familiar user, but this user may be a bit more demanding in her impression formation. As she already has an idea of what information she can find in this system, it may be more interesting for her to see what has happened since her last visit. This is not possible in the first generation tag cloud.

In the next chapter, we will show how a second generation tag cloud with more functionality, can benefit further from the strengths of the folksonomy. With these functionalities, we believe that the tag cloud can become a more efficient and user friendly browsing tool.

4 Second generation tag cloud functionality

In this chapter we will be looking at suggestions to functionality for a second generation tag cloud. This new functionality, will be examined in relation to how they may limit the previously stated limitations. We will show related work to each suggestion as well as discuss their contributions to a second generation tag cloud.

4.1 Structure in a second generation tag cloud

As stated earlier in this thesis, one of the limitations of the first generation tag cloud is that it does not show any relations between the tags in the cloud. This problem, as well as the fact that the alphabetical arrangement does not facilitate visual scanning and that the high semantic density of the most popular tags makes them bad discriminators, has been addressed by Montero et al in [9]. Montero points to these facts as limitations to the tag cloud as a visual information retrieval interface. To combat the problem of the high semantic density, Montero et al propose a method for the layout of a tag cloud where tags are grouped by similarity.

Earlier in this thesis we stated the problem that one or more topics tend to dominate the tag cloud with their tags. This problem has also been addressed by Begelman et al in [13]. They state that in a tag cloud it is hard to find one or two relevant tags because there are so many related tags scattered in the tag cloud. The fact that one or two popular topics and their related tags tend to dominate the tag cloud is a problem, because many interesting tags do not show themselves as they are drowned out by the more popular tags. In addition this popularity may lead to more use of the same tag [38].

To combat these problems Begelman et al proposes, in [13], an algorithm to cluster tags according to their co-occurences. In [9] Montero et al also use co-occurences analysis of tags to build clusters and visualize them in a clustered tag cloud.

By counting the number of co-occurrences, that is, the number of times when two tags are assigned to the same focus [9], Montero and Solana get a view of the similarity of different tags. They consider this tag similarity as a kind of semantic relationship between the tags. One can discuss whether or not this similarity is indeed a semantic similarity. The fact that two words have been assigned to describe the same focus does not necessarily infer semantic similarity. Montero and Solana base their relationship between co-occurrences and semantic similarity on the non-trivial nature of co-occurrence relationships described by Cattuto, Loreto and Pietronero in [39, 40] and the relationship between the tags measured by means of relative co-occurrence between tags [9]. This relative co-occurrence between tags is equal to the division between the number of resources in which tags co-occur, and the number of resources in which appear any one of two tags (the Jaccard measure). The equation can be seen in equation 4.1. They also call this co-occurrence-based similarity, tag overlapping.

$$RC(A, B) = \frac{|A \cap B|}{|A \cup B|} \tag{4.1}$$

By changing the way tags are selected for the tag cloud, Montero and Solana build a

Figure 8: A tag cloud before Montero and Solana

Figure 9: The same tag cloud after Montero and Solana

tag cloud where the overlap between tags are smaller. In addition Montero and Solana's result is a tag cloud where similar tags are grouped together in clusters and clusters are shown near similar clusters.

When compared to the original tag cloud in figure 8, one can see that their result in figure 9 shows that they have managed to remove some of the less significant tags (such as "toread", "diy" or "cool") and that some more interesting tags have emerged in the cloud (such as "philosophy" or "religion"). Several synonyms of the tags already in the tag cloud have disappeared and the clustering has limited the synonym problem inherited from the folksonomy. The clustered tag cloud is also no longer as dominated by technology-topics as it was before.

For additional readings on co-occurence analysis to infer similarity between tags in a collaborative tagging system, the reader is referred to a paper by Cattuto et al in [41], Michlmayr and Cayzer in [42] and Mika in [43].

The clustering of tags in a collaborative tagging system has also been proposed by Choy and Lui in [44]. Their proposal is to use latent semantic analysis (LSA) to discover similarity between tags. LSA is a technique in natural language processing for analysing relationships between documents, and the terms in them by producing a set of concepts related to the documents and terms. LSA is capable of eliminating noise, as well as extraction of latent relations in the data [44]. For more details on LSA, the reader is referred to [45, 46].

When stating the problem of synonyms, we mentioned that the use of both the plural and singular terms of a word is adding to this problem. In their studies Choy and Lui found that not all plural/singular tags (e.g. “movie” and “movies”) are related. They found that some plural/singular tags did not show any similarity at all (according to the LSA). In their LSA the range of similarity was -1 (highly dissimilar) to 1 (highly similar), and an LSA on the tags “photo” and “photos” found a similarity of 0.02. Other tags, such as for an example “recipe” and “recipes”, had a higher degree of similarity; 0.93. Choy and Lui found that in the case of some tags, the plural and singular terms of a tag carry references to specific genres.

The results of Choy and Lui show that similar tags, in a collaborative tagging system, may be easily located. And they found that by locating tags that did not fall in any specific categories, they could discover resources that had been tagged inconsistently.

Co-occurence and latent semantic analysis are two ways of inferring a structure on the tags in a collaborative tagging system. Another way is to discover a hierarchical structure as proposed by Brooks and Montanez in [47]. They argue that they can extract hierarchical information from a flat, non-hierarchical folksonomy and they are supported by Li et al in [34].

To infer this hierarchical relationship between tags, Li et al analyse the coverage of tags. This coverage is more closely defined in [34], but the overall theme is that a tag in a high semantic level is assigned to more resources than tags on a lower semantic level.

When observing the hierarchical structure from the semantic landscape Li et al notes that

> “[. . .] there is not a neat tree structure like taxonomy or human built ontology with rigid hierarchies and pre-defined categories with clear boundaries for social tags, but the tags used in social annotation do locate in different semantic levels in the social annotation space.”

They also list three features of this inferred hierarchical structure as

- There are multiple ways to the target resource. To find pages about JavaServer Pages we can make the clicks of “programming”, “java” and “jsp” in order or we can follow the path from “design”, “web” to “jsp”.
- The categories are not rigid. One URL can be associated with several categories, since different people have different views about the same thing.
- The hierarchical structure is dynamic with the increasing of social annotation data. If resources associated with a leaf tag increase, the leaf tag will extend several children nodes.

In figure 10 we show an example of Li’s hierarchical structure and how a leaf tag is a child of several categories. This feature makes it easier for a user to find a leaf node when browsing the hierarchy; if she chooses the “wrong” start-node she can still end up at the correct leaf-node. This is not always the case with traditional hierarchies [44].

When researching for this thesis, we were not able to find any commonly used tag clouds that implemented either clustering or an inferred hierarchy.

4.1.1 Discussion

By clustering tags in similar groups we can benefit from some of the benefits that hierarchical categorizing has, and we can try to avoid some of the limitations that the tag cloud inherits from the underlying folksonomy. Kuo et al found in [23] that tag clouds are inferior to lists when one wants to discover relations between concepts — And Mathes

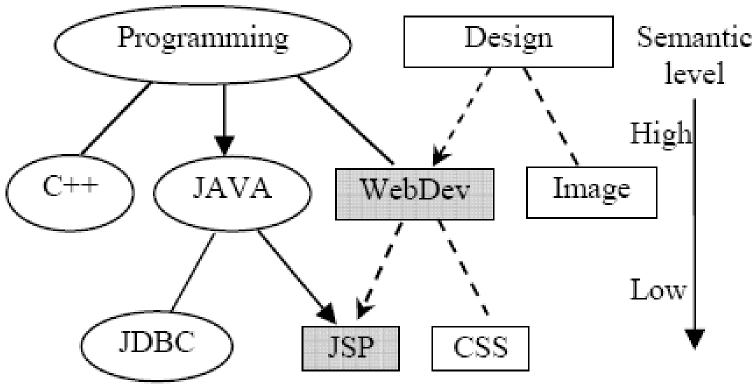

Figure 10: An example of Li et al's hierarchical structure [34]

described in [7] the problem of synonyms and polysemes in folksonomies. By comparing an unclustered tag cloud (figure 8), with a tag cloud that has been clustered by Montero and Solana (figure 9), we can see that some of the synonyms have disappeared and it is now easier to understand what some of the non-descriptive tags are about (such as the tag "design" is now clearly related to web design).

As we have shown, there are several suggestions as to how one can discover similarity between tags. This paper has not evaluated the effectiveness of either the Jaccard index or latent semantic analysis. But we do state that the effectiveness of the discovered clusters should be evaluated, to see how they perform in collaborative tagging systems.

It is the opinion of this thesis that the new and clustered tag cloud facilitates visual browsing and scanning a lot better than the tag cloud that was arranged alphabetically.

Then again, Halvey and Keane discovered in [21] that alphabetization can be helpful to the end-users of a tag cloud. Of course to properly evaluate if a clustered tag cloud is more useful it will be useful to perform user studies. Most likely the varying needs of the end-users are what will dictate whether or not the tag cloud should be clustered or not, which again strengthens the needs for the ability to manipulate the presentation of the tag cloud so that the end-user can choose for herself.

4.2 Time utilization in a second generation tag cloud

A temporal tag cloud is a tag cloud that implements utilization of time to allow the user to choose the time frame for the tag cloud. We have stated earlier that one limitation of the first generation tag cloud is its static nature and inability to visualize the evolution in a collaborative tagging system. We stated that a familiar user could benefit from the ability to monitor the change that has occurred in the collaborative tagging system, since her last visit. This section will show related work on visualizing tags over time, and why we believe there is a need for it in the tag cloud.

A tag cloud provides its users with a snap-shot of the most used tags in a given time frame. Sometimes this time frame is a day, a week or even all time. The first generation tag cloud is very static, and does not visualize the development in the tag space over time. This problem limits our ability to discover trends and see what topics interests the community behind the tag cloud. The web page Technorati [3] is an Internet search engine for web logs, and as of June 2007 it was tracking 86.8 million web logs [48].

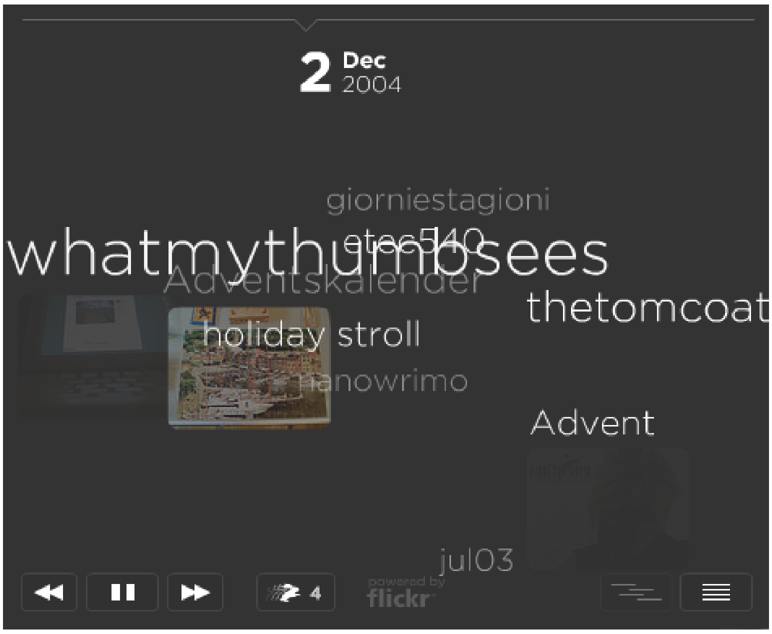

Figure 11: A picture of the Tag Lines interface

With a temporal tag cloud we can monitor what this web log community is writing about at a given time.

During the Web 2.0 Conference in 2005 Tim O'Reilly said that by watching the emergence of the tag "Ajax" they could predict that if O'Reilly published books about Ajax, readers would gobble them up [49].

In 2006, Dubinko et al [50] developed Tag Lines [51]. This project visualizes the evolution of what Dubinko defines as "interesting tags" from Flickr. In Tag Lines a user is presented a visualization where the time gradually shifts from June 2004 and through to September 2005. As time passes, so does the tags and the user can at any time click on a tag to view the photos that are associated with that tag. The user is also able to manipulate the time by dragging a slider at the top of the presentation. By giving the user different controls to manipulate the presentation (such as pause, fast forward or backward), Tag Lines has found a way to present the evolution of the semantic information in an interactive way. Figure 11 shows a picture of the Tag Lines interface.

A different way to visualize the evolution of tags in a tag space is shown on Cloudalicious [52] by Terrell Russell [53]. Here a user can type in the url of a web page and she is then shown a graph of which tags this web page has been tagged with on Delicious. Cloudalicous shows how taggers have conceived a web page over a period of time and can be interpreted as a temporal tag cloud visualized as a graph. This graph adds one more piece of information to the tag cloud; the evolution of the tag cloud. A picture of a Cloudalicious graph of CNN.com can be seen in figure 12.

The temporal aspect of the tags in a tag cloud does not necessarily have to be development over time. In [23] Kuo et al describe their system PubCloud that visualizes search results from queries over the PubMed database of biomedical literature. The visualization presented by PubCloud is equal to the visualization of a tag cloud apart from the visualization of recency of tags. To visualize the recency of the tags, PubCloud colours

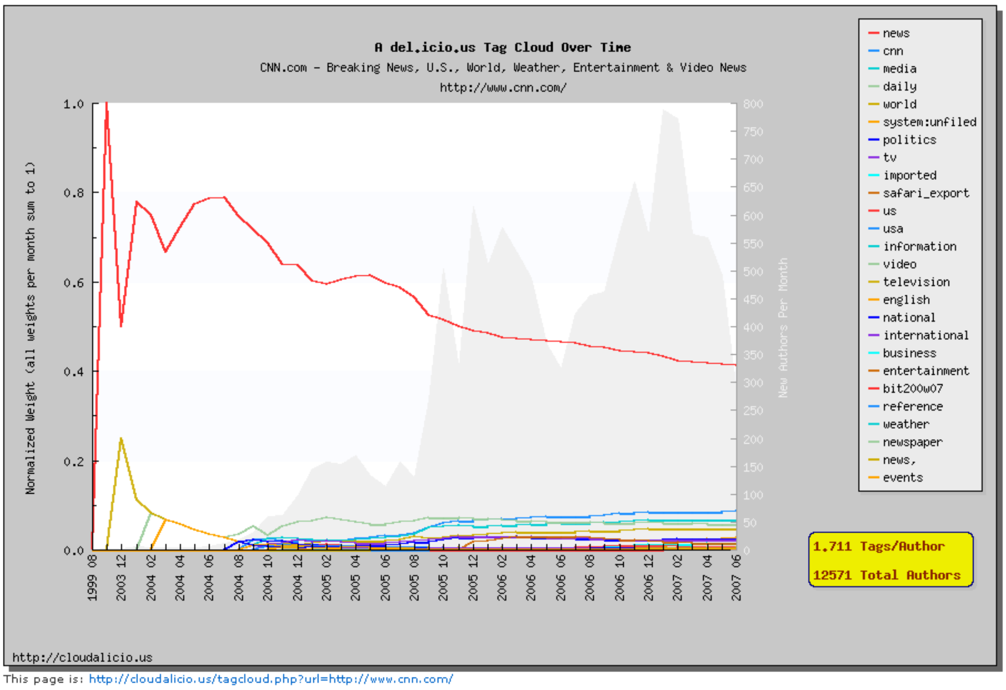

Figure 12: A cloudalicius graph of www.cnn.com

the tags in the cloud ranging from bright red for the most recent to dark grey for the oldest.

The static nature of the first generation tag cloud fails in visualizing important information to its viewer. This problem is often emphasized in tag clouds that visualize a long period of time. Some services have a tag cloud that visualizes the most popular tags of all time — an example of this is Flickr's "All time most popular tags" shown in figure 5. In a tag cloud like this one there is very little change and a new tag will have difficulty being noticed among the others who have had a long time to gain their popularity. In the tag cloud in figure 5, we can not say for sure that the tag "japan" is popular at the moment, only that it has been very popular before.

As the tag cloud emphasizes the most popular tags in a tag space, less popular tags have a tendency to be drowned out by the more popular or more used tags. The rate at which something, in this case a tag, loses its popularity is called popularity decay [14]. In the previously mentioned tag cloud from Flickr the popularity decay is never reset and for a tag to become visible in this cloud it would have to become more popular than the other tags in the cloud which have had a very long time to gain popularity.

This brings us to the topic of emergent tags. Emergent tags are tags that are in the process of becoming popular. These tags are not among the most used but they are gaining popularity fast. Emergent tags are valuable because they show trends and the development of trends [49]. By looking at emergent tags from forums, blogospheres etc. one can catch upcoming hypes, trends and the like while they are still in the beginning.

As stated by Rech in [54]

> "Trends are a valuable tool to shape one's research, decide on acquisitions, inform about technologies or standards used in the development process."

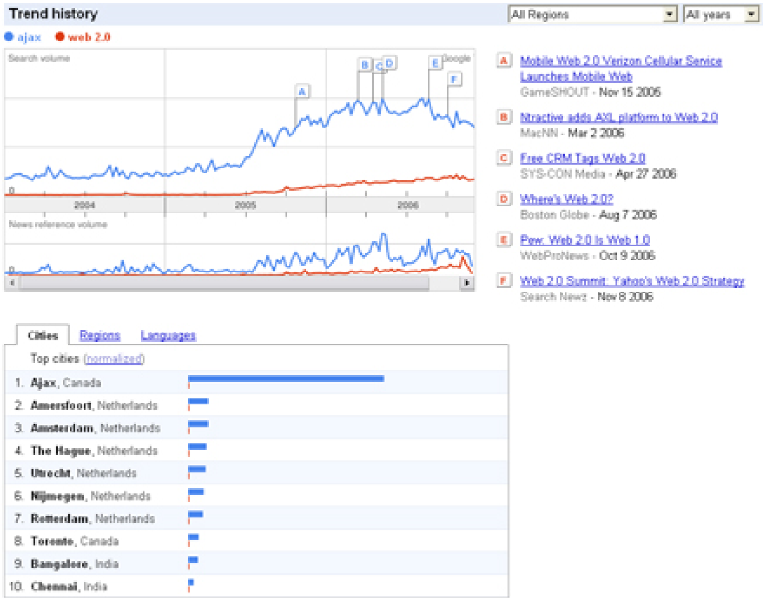

Figure 13: A result from Google Trends

Popularity decay and emergent tags are closely related as one needs an evolving tag cloud to monitor the emergent tags. The earlier mentioned tag cloud from Flickr is not a good way of detecting emergent tags. To try and combat this, Flickr has three stages of tag clouds. The first stage displays the most popular tags the last 24 hours. This stage has a short popularity decay i.e. a tag that is popular today may not be in the cloud tomorrow. The second stage displays the most popular tags the last week and the third stage is the previously mentioned "All time most popular tags". By displaying these three stages of popularity decay, Flickr is giving their users the opportunity to get more of a temporal aspect to the tag cloud and the information it is displaying.

The need to use information visualization to discover emerging trends and patterns over time is not new. In 1996 Nowell et al [55] developed a system called Envision that visualized search results in a digital library of computer science. They found that when asked, their interviewees would most like to have the ability to identify and explore patterns in their search result.

The visualization of trends and evolution of interests has also been addressed by Jörg Rech in [54]. Rech describes a service called Google Trends [56]. Google Trends is a service provided by Google which can be used to discover trends — this service was opened to the public in May 2006. By entering up to five terms (or phrases) a user is presented with a graphical view of the search-, and news-history of these terms. The results are based on searches to Google Search and news articles collected in Google News [54]. The result of a search for "ajax, web 2.0" is shown in figure 13, the top graph shows search volume and the lower graph shows news reference volume. The user is allowed to interact with the result by choosing which country the result should be chosen from, and over which time line the result should be based upon.

As far as we have managed to gather, there are no implementations of temporal tag clouds in use today. And the only research we have managed to find that visualizes tags over time is Tag Lines.

4.2.1 Discussion

This thesis believes that one of the greatest limitations of the tag cloud is its failure to visualize temporal changes in the tag space. The tag cloud was firstly and foremost implemented as a way to visualize which tags are most popular at a given moment. But, popularity shifts and something that is popular today may not be popular in a week or two. This information is not visualized in the first generation tag cloud and we consider this to be a serious drawback. With a static first generation tag cloud and a long time frame (e.g. Flickr's "all time most popular" tag cloud), we are only provided with a view of which tags have been popular in the system; we have no way of knowing if these tags are still popular or if they are no longer in use.

Earlier we mentioned Flickr's three stages of tag clouds. The benefits of these different stages are, as mentioned, that it is easier for emergent tags to be noticed. However these tag clouds are not what this thesis considers as tag clouds that show evolution over time. The three different tag clouds at Flickr [57] do not show us how the user's tagging preferences have changed e.g. from day to day in the last week. We would like to have the ability to closer study the evolution over time e.g. by comparing the different tag clouds of successive days or weeks. Either by a way of dynamically changing the time frame of the tag cloud or by comparing different tag clouds side by side.

A way to dynamically change the time frame of a tag cloud could be to implement a slider. By dragging this slider the time frame could be changed and the tag cloud would change with the time. The user could then see how some tags grow in size and others diminish. With regards to discovering emerging tags and monitoring the evolution of a collaborative tagging system, we state that one dynamic tag cloud is better than three static ones.

Some systems do visualize temporal change both in tagging systems and elsewhere. One of these systems is Cloudalicious [52] by Terrell Russel [53]. This visualization is however not what this thesis envision as a the temporal information to be visualized in a tag cloud. The information relayed by Cloudalicious shows us how different taggers on Delicious have perceived one specific web page. As we described in the introduction to this thesis, the "gist" a focus (or resource), can be found by studying the tag cloud for that focus. This is what cloudalicious does. For a second generation tag cloud, we would like to see all the tags that have been assigned to all the resources, not just one web page.

The web service Google Trends [56] is described by Jörg Rech in [54]. This kind of visualization of trends is more along the lines of what we envision as an implementation to tag clouds. By changing the query to something like "what terms are people searching for" or "what terms are web loggers writing about" and changing the visualization to a tag cloud, one could end up with a tag cloud that visualizes trends in e.g. a web log community like Technorati.

The system closest to what this thesis envisions as a temporal tag cloud is Taglines [51] by Dubinko et al [50]. The visualization in Taglines is interesting because it shows us that it is possible to visualize the evolution in a tag space in an interactive way, but it should be transferable to other sites and other resources than photos. An alternative to the approach in Tag Lines is a tag cloud that more or less looks like a first generation tag cloud but with the option for the user to change the time line (e.g. by dragging a slider) and watch the tag cloud change in real time.

Inside a semantic field, such as a tag space, there is so much more information than just which tags have the highest frequency of use. If the tagging systems were to assign a time stamp to each tag entered into the system, one could easily build a temporal tag cloud that shifts its time frame as the users commands it to. But why only visualize the most popular tags? With time stamps the system could also be instructed to visualize which tags are accelerating, in popularity, the most. This way a user could be presented with a cloud of the tags that have been added most often during a specified time frame i.e the tags have had the highest increase in usage during e.g. the last week. This information could probably be visualized with the help of colours as proposed by Kuo et al [23].

4.3 Presenting a second generation tag cloud

In the introduction to this thesis, we presented some limitations to the first generation tag cloud. The last three limitations that we listed, were the lack of a detailed view and lack of interactivity, as well as the problem that one word is a very wide search-field. This chapter will be aimed at showing related work on these issues. We will, in this chapter, show how the presentation of the tag cloud can affect it as a visual information interface. We will also look at how the wide search-field of only one word can be narrowed down to a more specific search-field. And we will explain why we believe that the user should be able to get a more detailed view of the tag cloud.

4.3.1 The layout of the tag cloud

In [21] Halvey and Keane present an assessment of tag presentation techniques in tag clouds. By evaluating different properties that can be utilized in presenting tags (such as font size, alphabetization etc.) they investigate the effect that these properties have on the ease with which users can find tags. Their study shows that alphabetization can aid users to find information more easily and quickly, and that both the font size and the position of the tags in the tag cloud is very important to how easy and quickly users find information. Their experiment used goal-oriented task as opposed to a free-browsing task. In our thesis, we focus mostly on the tag cloud as a browsing interface, but their research is still interesting.

The importance of the position of tags in the tag cloud is also proposed by Rivadeneira et al in [10]. They propose that future designers of the tag cloud, should focus on the layout and that the designer can rely on established psychophysical principles in setting basic tag cloud parameters. Rivadeneira et al states that these parameters produced results predicted by perception theory, namely that font size and location affected (low-level) memory. However it may not be as simple as to state that large fonts are noticed more easily; Halvey et al discovered in [21] that some participants found big tags to be hard to find.

In their development of PubCloud in [23] Kuo et al also notes the need to implement interface interactivity to allow the user to change colours, font sizes etc. In addition to the ability to choose whether or not to cluster the tags in the presented cloud.

Some systems have already implemented the ability for their users to change some of the layout of the tag cloud. The bookmarking site Delicious implements a functionality called Tag Rolls [58], where the user can build her own customized tag cloud by setting different parameters — such as font size, font colour, arrangement of the tags and size of the cloud — for the visualization. In figure 15 we show a picture of the display options

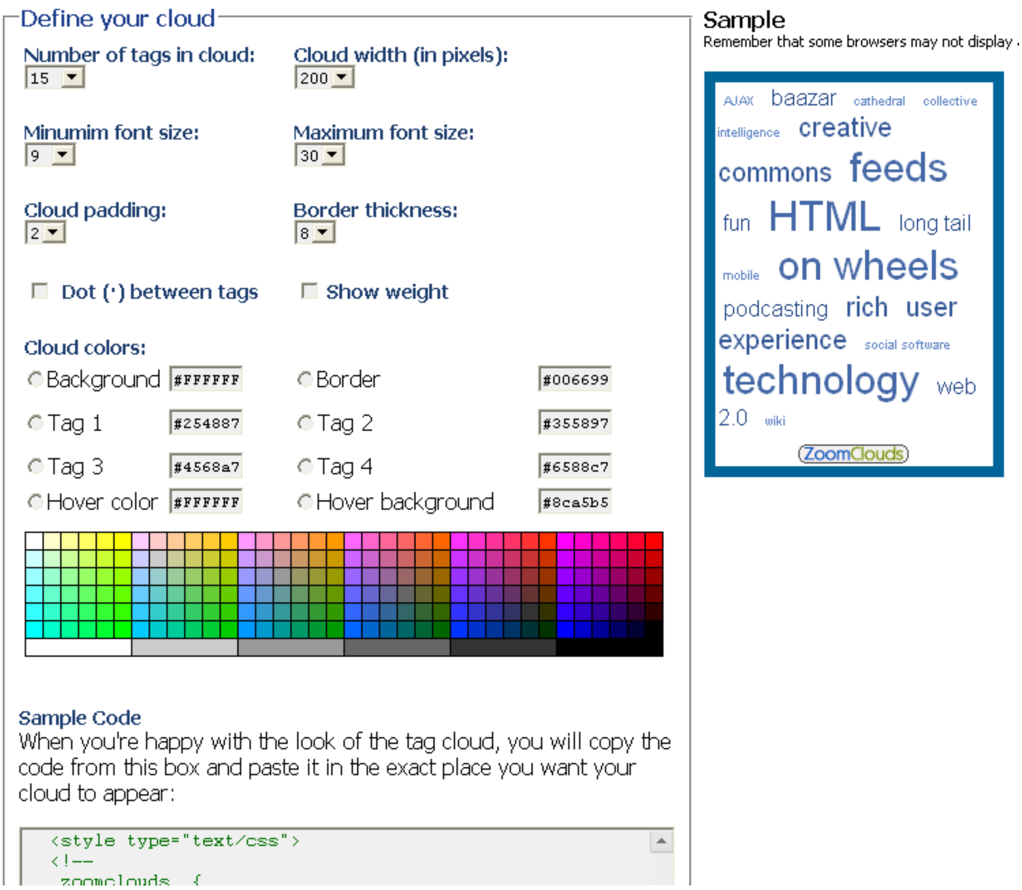

Figure 14: The interface from ZoomClouds

on Tag Rolls.

Another service that allows a user to define a tag cloud is ZoomClouds [59]. ZoomClouds is a website that, using the Yahoo! Content Analysis API, produces tag clouds for any given RSS-feed. When the user has chosen an RSS-feed to base the tag cloud on, she is given a number of options to manipulate the presentation. A picture of this interface is shown in figure 14. The service produces CSS- and HTML-code which the user can then paste on her website.

These two interfaces (Tag Rolls and ZoomClouds) are meant solely for manipulating the tag cloud before posting it, and not for dynamically changing the presentation of the tag cloud once it has been posted to a website.

In [60], Ahlberg and Shneiderman describe important principles when designing systems for visual information seeking (VIS). The principles we found most interesting to our work are listed below

- Rapid, incremental and reversible actions
- Selection by pointing (not by typing)
- Immediate and continuous display of results
- Tight coupling

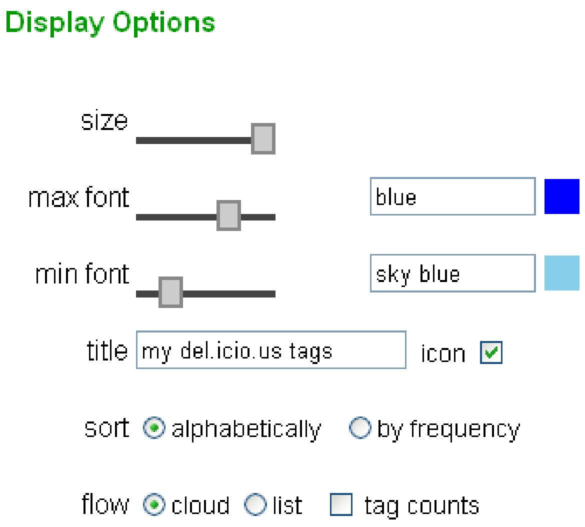

Figure 15: The display options on Tag Rolls from Delicious

Tight coupling is described by Ahlberg and Shneiderman as the fact that query components are interrelated and reflect the changes made with them. An example of tight coupling is the save-button on many text editors. As the document is saved, the button becomes grayed out until a new change is made in the document. Another example is the handle on the scroll bar on a web page; as the user scrolls up or down on the page, the handle is moved on the slider accordingly.

The first generation tag cloud is ordered as lines of text where each line contains a number of words. Because these lines contains large and small text interspersed, the result is often a tag cloud with a lot of wasteful white space. This problem has been addressed by Kaster and Lemire in [61]. They propose to use electronic design automation (EDA) to automatically place the tags in the cloud. By doing so they hope to construct a more aesthetically pleasing tag cloud. And a tag cloud that is easier to view on small screens, such as on a PDA.

The standard line-based layout of the tag cloud, is not the only way to visualize a tag cloud. In [62] Bielenberg and Zacher propose a circular layout for the tag cloud. In their layout the font size of the tag in the cloud and its distance to the center of the circle, represent the importance of the tag (an example can be seen in figure 16).

In [63] Blake Shaw show us a presentation where he visualizes the tags as a graph where each tag is a node and the similarity between each tag is an edge between the nodes.

This thesis will not be concerned with other ways of visualizing the tag cloud than the previously mentioned line-based layout. It is however interesting to note that there has been done research on visualizing information in two and three dimensions. Westerman and Cribbin state in [64], that the third dimension does not appear to add a significant advantage over two-dimensional displays. Westerman and Cribbin found that the amount of additional semantic information that can be conveyed by a three-dimensional solution does not outweigh the associated additional cognitive demands. This discovery is supported by Sutcliffe and Patel in [65].

As far as this thesis has managed to find, the closest implementation to a tag cloud

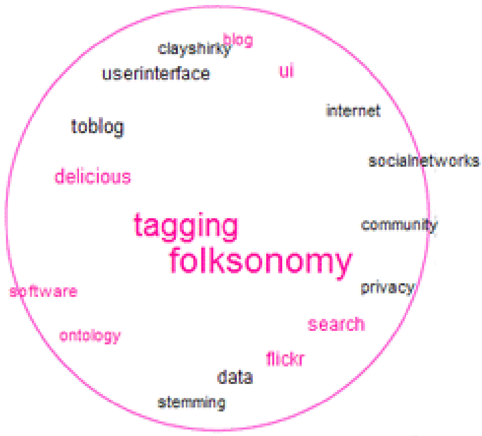

Figure 16: A circular tag cloud by Bielenberg and Zacher [62]

that allows its users to manipulate it, are the services Tag Rolls and ZoomClouds. But as mentioned, they only manipulate the tag cloud before posting; The reader of the tag cloud can not change the presentation at a later stage. We have not been able to find any tag clouds that allow the user to dynamically change the presentation on the tag cloud she is currently reading.

4.3.2 A detailed view

In the introduction to this thesis we mentioned Shneiderman's visual-information-seeking-mantra, "overview first, zoom and filter, then details on demand". And in [60] Ahlberg and Shneiderman describe the ability to zoom in on a result set to reduce cluttering and get a more detailed view. When researching for this thesis we have not been able to find any implementations of the first generation tag cloud that allow for a more detailed view. Later in this chapter we will show how we have developed our proof of concept to allow the user to zoom in on the tag cloud that she is viewing. Our zooming-functionality allows the user to "dive" into the tag cloud. It is much like zooming in on a satellite image; the largest tags disappear from view while the others increase in size. By doing so we hope to give the user an ability to look beyond the most protruding tags, and thereby getting more information than just the standard "which tags are the most popular".

4.3.3 A narrowing search-field

Previously we mentioned Ahlberg and Shneiderman's principles for the design of visual information seeking systems [60]. One of the principles points to incremental actions. Ahlberg and Shneiderman also mention the design of systems where the query of the user is progressively refined.

In a tag cloud, the search query is the tag that the user chose in the tag cloud. This search query is very wide, and we propose that the user could benefit from the ability to refine her search query. Faceted browsing describes a way the user can do so. Both Sinha [66] and Yee et al [67] describes faceted browsing. In our description on folksonomies, we explained facets as different aspects of a resource. Faceted browsing can be used in the tag cloud interface to allow the user to refine her search-field.

In faceted browsing, when a user clicks a tag in the cloud, instead of being presented with the related resources, she is presented with a new tag cloud. This new tag cloud contains the tags most used in conjuncture with the previously chosen tag — e.g "show me the tags that are most used in union with the tag 'ajax' ". Some systems already implement faceted browsing, two of these are Wine.com [68] and Rawsugar [69].

4.3.4 Discussion

With the first generation tag cloud there is generally one presentation technique; the standard line-based layout. The tags in the cloud may be ordered alphabetically or by frequency (as seen in Tag Rolls). As the possibilities for presenting the tag cloud increase, we believe that the user should gain the opportunity to choose how she wants her tag cloud presented. We see the tag cloud as a user interface that should change according to the user's wishes. Not a static information visualization that the user can not manipulate.

As of today, the lack of alternative presentation techniques naturally leads to a limit in possibilities in manipulating the presentation of the tag cloud. But we have shown that even small changes like the font sizes or the position of the tags in the cloud, can affect the usability of the tag cloud.

By allowing the user to choose whether or not she wants her tag cloud to be clustered, choose the time frame etc., we hope to give the user a greater impact on the information that is visualized to her. With a tag cloud that changes at the whim of the user, we hope to give the user a greater impact on the information that is visualized to her.

By referring to Shneiderman's visual-information-seeking-mantra, we have stated that a tag cloud user may benefit from the ability to zoom in on a tag cloud. And in the next chapter, we will show how we have implemented this in our proof of concept.

We have also stated that the search query (i.e. the chosen tag) from a tag cloud, is very wide. With a very wide search query, the result set will also be very wide. For a user who is searching for information, this wide result set will most likely contain many unrelated resources. Adam Mathes describes in [7], how his search for "filtering" returned a result set containing resources on both Bayesian filtering and the filtering of vodka. With a tag cloud implemented with functionality for faceted browsing, Mathes could instead have been presented with a tag cloud of words used together with "filtering". This way the user will not have to search through a potentially large result set to find a related resource. The opinion of this thesis is that faceted browsing can make the tag cloud a more effective and less frustrating interface for information seeking.

4.4 Summary

In this chapter we have shown some suggestions to functionality in a second generation tag cloud. These suggestions have been discussed as to how they can limit the problems of the first generation tag cloud, which we listed in the introduction to this thesis.

In the first section, we showed that even though the the structure of a folksonomy is flat and non-structured, it is possible to infer both a similarity between tags, and a hierarchy among them. By visualizing this structure in a tag cloud, the problem of lacking structure and missing relations in the first generation tag cloud, can be limited. Earlier in the thesis we mentioned that browsing is a superior strategy when "there is a good underlying structure so that items close to one another can be inferred to be similar" [25]. As we saw in the clustered tag cloud by Montero and Solana (figure 9) they have also, to some extent, limited the problem of dominating topics in the first generation tag cloud.

With the basis in these points, we propose that Montero and Solana's tag cloud has increased the first generation tag cloud's usability as a browsing interface.

We have shown that co-occurrence and latent semantic analysis are two methods to discover clusters in the tag cloud. For future work we could like to see research that evaluates just how similar the tags in the clusters are, semantically speaking.

This thesis propose that by clustering the tags in the tag cloud, we can improve on the benefits it already has as a browsing interface. Specifically we believe that the user-roles we have defined as "unfamiliar", "familiar" and "browsing", may benefit from this clustering. We base this on the belief that it is easier to gather information from a clustered tag cloud. We propose that it is easier to understand what the different tags are about in a clustered tag cloud. With the help of the clustered tag cloud a user can gather that the tag "ajax" is related to XML and JavaScript, and not the Greek hero or the Dutch football club.

In the introduction to this thesis we stated that the lack of time control was one of the limitations of the first generation tag cloud. In the second section of this chapter, we showed why we propose to implement time in a tag cloud. We have shown why the ability to manipulate time can be a useful functionality in a second generation tag cloud. With time-utilization in a tag cloud, we state that the users can discover and monitor changes in the underlying collaborative tagging community. With this in mind, we propose that the user-role we defined as "familiar", may benefit from such a utilization. When researching for this thesis, we did not manage to find any tag clouds with implemented time-utilization.

The last three limitations, of the first generation tag cloud, which we listed in the introduction where:

- No detailed view,
- Very wide search-field and
- No interactivity.

Through this chapter we have stated that the tag cloud can benefit from a more detailed view and the ability to limit the search-field. We described faceted browsing, and how it can allow the user to narrow her search-field. We have not been able to find tag clouds that implement any functionality to get a more detailed view.

In this chapter we have shown how the layout of the tag cloud can affect its usability. And we have shown examples of systems that allow the user to manipulate some of the presentation of the tag cloud. If the functionalities shown in this chapter are implemented in a tag cloud, we state that this tag cloud will need the option to change the way it is presented. We propose a design for a second generation tag cloud where the user can choose which functionality to use. An example of this, is a tag cloud where the user can choose whether or not to cluster the tags in her tag cloud.

To evaluate these new functionalities, we will relate them to our previously defined user-roles:

- Unfamiliar user
- Familiar user
- Browsing user

- Searching user

Firstly, we will be looking at the inferred structures in the second generation tag cloud. With regards to the two first user-roles, they are now dealing with a more "tidy" tag cloud. As we have shown, some synonyms of the dominating topics disappear, and so has some less significant tags. By reducing the domination of the most popular topics, it becomes easier for any kind of user to find interesting tags. And the less significant tags ("toread", "diy") are not of much use by their own. We have previously described how an unfamiliar user can use the tag cloud to form an impression to get a sense of the underlying information. With a clustered tag cloud, it is easier to understand what the different tags are about. And we state that this should help the unfamiliar user to form her impression. The same goes for the familiar user; a more comprehensive tag cloud is beneficial for everybody. We state that for both the familiar and unfamiliar user, the clustered tag cloud relays more useful information than the unclustered tag cloud.

We have previously explained how the browsing user can benefit from the first generation tag cloud. And according to Lin [25], browsing is a superior strategy when "there is a good underlying structure so that items close to one another can be inferred to be similar". In the first generation tag cloud, this point was not satisfied. In the clustered tag cloud, the point is satisfied. By fulfilling one more of Lin's points, we state that the tag cloud has become a better browsing tool. Of course, user testing could help find whether or not this is actually true.

When looking at the searching user and the first generation tag cloud, we found that it seems that the tag cloud is not a good tool when *searching* for information. This may, amongst other things, be due to the lack of structure. We compared it to looking for a specific book in a library where the books have been placed randomly around in the shelves. With more of a structure in the tag cloud, be it either clusters or an inferred hierarchy, we state that the finding user will have an easier time finding her information. The tag cloud still seems to be inferior to a search engine, but with a structure, it may be easier to find specific information.

Secondly, we will be looking the time-utilization in the second generation tag cloud. In our explanation of the user-roles, we stated that a familiar user already knows the tag cloud and the underlying system. This user may be in need of more information than just the impression formation of the unfamiliar user. As we previously explained, in the first generation tag cloud there is no possibility for the familiar user to see how the tag cloud has changed since her last visit. With the utilization of time, she can now find these changes. The unfamiliar user can on her side, use the time-utilization to further improve her understanding of the underlying information. By manipulating the time frame, she can gather information on how the interests of the community has shifted over time.

When defining our user-roles, we explained the browsing user as one that does not have a clear question that needs to be answered, but she has an idea of what kind of information she is looking for. We do not see this time-utilization as helpful for the browsing user, unless she is browsing for information on how the interests of the community has evolved. So far we have based our evaluation of the browsing user on Lin's list. In this list, there is nothing that indicates that browsing can benefit from time utilization.

The searching user does not seem to have much use for the tag cloud, so far. And the way we see it, the time-utilization has not improved on this. When searching for information related to a specific question, the time-utilization does not seem to help the searching user in getting a more specific result-set. The wide one-word-search-field of the first generation tag cloud is still present.

Lastly, we will look at how faceted browsing and zooming-functionality will affect our user-roles. With regards to both the unfamiliar user and her impression formation, faceted browsing does not seem to provide her with any useful functionality. We state that the impression is formatted when the user looks at the tag cloud, and not when she clicks on a tag to find a resource. The zooming-functionality can however provide her with more information on which tags are not as popular. When the unfamiliar user looks at the tag cloud and the largest tags, she can quickly see which words are most used. By zooming in on the tag cloud, she can also discover words that are in use, but not with the same frequency as the more popular ones. With the familiar user, the situation is a bit different. As she is already familiar with the popular topics in the system, she might be interested in having a look at which topics are not as popular. Tag clouds visualize popularity, and if this popularity does not shift, the largest tags will stay the same. By zooming in on the tag cloud, and excluding the largest tags, a familiar user can gain new information about the community that she did not previously have. As with the unfamiliar user, we do not see that faceted browsing has any impact on the familiar user.

The searching user has previously been stated to be impaired by the tag cloud and its one-word-search-field. This changes with faceted browsing. With faceted browsing, the user now has the possibility to refine her search-field to be more than just one word. As with the browsing user, she also can refine her search-field so that her browsing is no longer confined to finding one related word. This should indicate that a tag cloud with faceted browsing capabilities, is more useful to both the browsing and the searching user. With regards to the zooming-capabilities, neither the searching or the browsing user seem to gain any benefits, other than the ones stated for the unfamiliar and familiar user (the ability to discover less popular tags).

As explained, we did not manage to find any commonly used tag clouds with time-utilization or zooming capabilities. In the next chapter, we will show how our proof of concept has implemented these functionalities.

5 A second generation prototype

When researching new functionalities for the second generation tag cloud, it became clear that we wanted to develop a system that implemented some of these functionalities, in a proof of concept. In this chapter we describe the proof of concept that has been developed. This proof of concept has been developed to show how a second generation tag cloud may look and behave.

In this chapter, we will describe the overall design of the system. And we will describe the functionalities that we have implemented, and how this has been done. At the end we will analyse our new second generation tag cloud, with regards to its new functionality and which consequences their implementations have had on the tag cloud. To do so, we will again be looking at our pre-defined user-roles:

- Unfamiliar user
- Familiar user
- Browsing user and
- Finding user

5.1 Overall system design

The proof of concept consists of a MySQL [70] database of 34 news articles from the news archive of Gjøvik University College [71]. Each article has been tagged with a minimum of 2 tags. As a result the database consists of 133 news-tags-relations. The back-end of the system is written in the scripting language PHP [72] and the front-end of the system is written in a combination of HTML, CSS and PHP. The front-end refers to the graphical user interface (GUI) that the user is presented in her web browser. The back-end is a PHP-script that receives variables and information from the front-end, and returns information to be presented to the user.

The design of the database can be seen in figure 18. And will be described more closely in the next section. Our test-station was a PC running Windows XP, service pack 2, with an AMD 2.21 GHz processor and 2GB of RAM.

After implementing new functionalities in our proof of concept, we wanted to compare this new second generation tag cloud with the first generation. We will be looking at new demands on the back-end (the database and the system that gathers data from the database) and the front end (the graphical user interface). For the comparison of the two systems, we will be looking specifically at the differences in functionality and complexity. When writing this thesis, we have not been able to find any specifications on how a collaborative tagging system (CTS) is designed. So to compare our system with a standard CTS, based on the minimum requirements needed to build a CTS.

In figure 17 we show a picture of the graphical user interface of our proof of concept, and the controls we have implemented in it. The different parts of the interface are detailed below.

A This is where the user can enter a minimum and a maximum frequency for the tags to

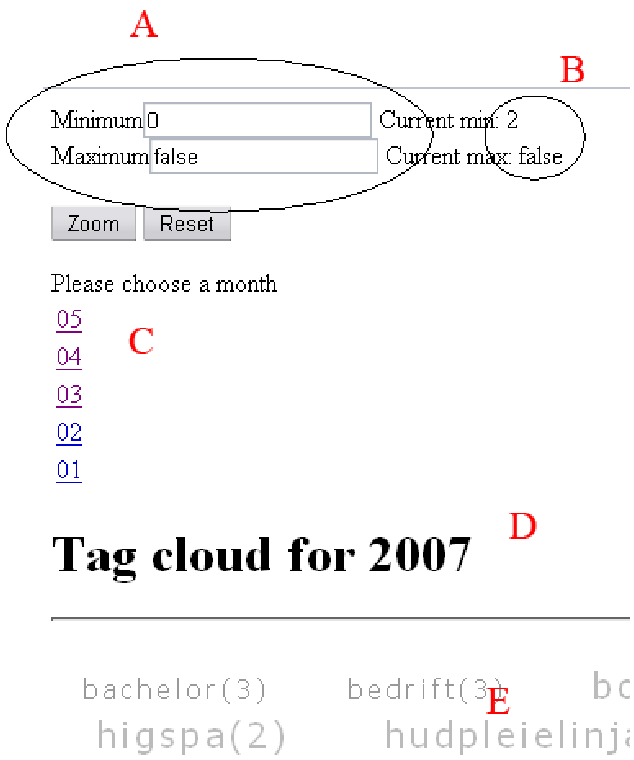

Figure 17: The graphical user interface of our proof of concept

be shown in the cloud. For an example, if the user entered "2" in the minimum-field, the tag cloud would not show tags with a frequency lower than 2. The opposite applies to the maximum-field.

B These fields show the user the current zoom-range she has chosen.

C This part of the GUI allows the user to set a time frame for the tag cloud. As the system is loaded the time frame is empty and the user is shown a set of years to choose from (these years have been taken from the database). If the user chooses a year, she is presented with the tag cloud for that year and her choices of time frames are changed to a set of months (as seen in figure 17). The tag cloud is then changed to present the tags for the chosen month. As the user has chosen a month, she is presented with a set of dates from that month and if she chooses a date, the tag cloud for that specific date will be presented to her.

D Here the current time frame is presented to allow the user to always keep track of the time frame for the current tag cloud.

E This is where the tag cloud is presented.

5.2 Time functionality

In our proof of concept the implementation of time-utilization was considered very important. This section will describe how this control was implemented and how the system behind could be designed to facilitate the time-utilization.

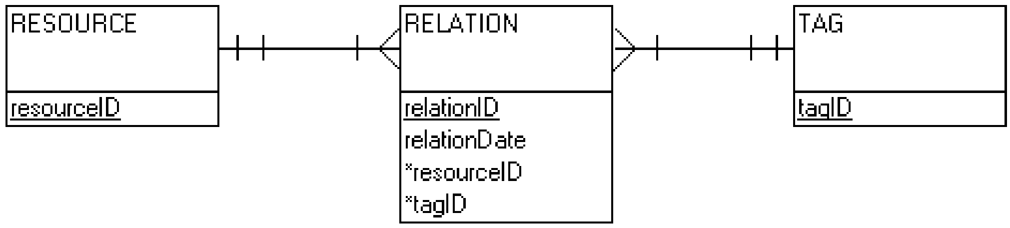

Figure 18: A model of the database design

The proof of concept is built upon a database which holds each resource and the tags that have been used to describe any one resource. To enable the system to differ between the time that each tag has been assigned and to which resource, the database consists of three tables. A model of these three tables is shown in figure 18. The table *resource* holds each resource that has been added to the system, and the necessary information about it (e.g. "URL", "title" etc.). The table *tag* holds each tag that has been assigned to a resource. The last table *relation* holds one entry for each tag-resource-relation.

When a tag is assigned to a resource, the id of the resource and the tag is added together with the date the relation has been created. As we can see, the id for the table *relation* is the attribute "relationID". This way resource "a" can be tagged with tag "b" several times on the same day without creating an error. By keeping track of the date in the *relation* table, the system can show us which tags have been assigned to which resources on a specific date (if one would want to examine the development even, closer one could add a "time" attribute to the *relation* table).

In figure 19 we show a flow chart-diagram of the part of the system that lets the user set the time frame. The first time the system starts, it gathers the years that have been used and lists these to the user. When the user chooses a year, the system builds and visualizes, a tag cloud for that year and lists the months that are available within the chosen year. The last choice for the user is to choose a day within the chosen month. Each time the user chooses a year, month or day, the activity is started over with the chosen time frame.

5.2.1 Analysis

The first part of the collaborative tagging system (CTS) we want to examine closer is the back-end. By the back-end we refer to the database and the system that collects data from the database. In our database design (figure 18) we have three tables that contain the resources, the tags and the relations between them. The table *relation* contains the tagged resource, the tag that describes the resource and the date the relation was created. The minimum requirements in the database of a standard CTS, we have defined as the same three tables; because of the many-to-many-relation between resource and tag. However, the date-attribute in the relation-table is not required in a standard CTS. In our proof of concept, this date-attribute is required so that we can log the date that each relation is created. When comparing these two database designs, we can see that there are not any difficult new implementations in our database design. However, if one is to upgrade a standard CTS (without the date-attribute) to our database design, one will face the problem of the missing attribute "date" (as the dates of the previously created relations are not logged).

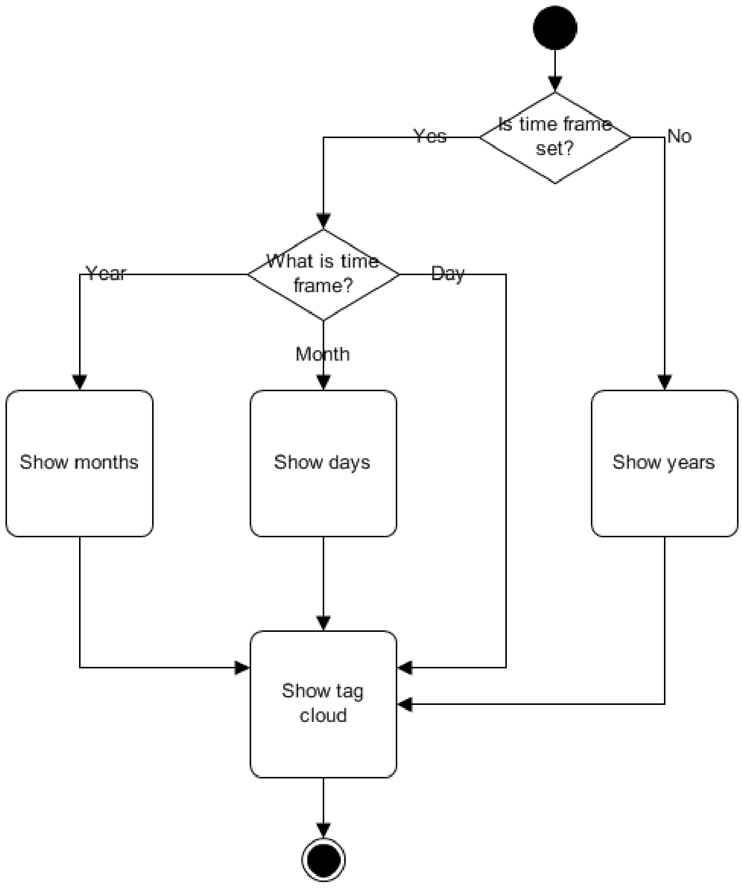

Figure 19: A flow chart diagram of the time implementation

The next part of the back-end we want to compare is the system that receives the tags and the resources when a user is uploading a new resource to the system. In our definition of a standard CTS, the system logs the date of the first upload of the resource. However, if a user decides to change the tags, or add tags to the resource at a later time, the standard CTS will not log the date of this change. In our system, we use the date-attribute both as the date of the first upload and the date of the tag-changes. The only difference between the two systems, is that the system will have to place the date in a different table at the first upload. And if tags are added to a resource at a later stage, the system will have to log the date that this happened. We do not view these changes as difficult to implement as they result in very small programming changes in the system.

So far we have covered the part of the back-end that receives a resource when it is uploaded and the part that handles changes to the tags on a resource. We will now take a look at the part of the back-end that collects the tags needed for the tag cloud. In a first generation tag cloud, tags are collected on the basis of their frequency of use. In our system we also collect the tags on the basis of the date they where added. The changes needed to accomplish this are minimal; all the developers will have to do, is change their SQL-statements in the system that collects the tags.

On the front-end of the system, there also needs to be made some changes in order to implement our time-utilization. When developing an interface that allows the user to choose a time frame, we believe it is important to consider the complexity of the new interface. We have previously stated that the simplicity of tagging and the tag cloud is what has made it so popular. In our quest for a second generation tag cloud, we had to keep this in mind so that we did not remove this simplicity.

When a user is uploading a resource or changing the tags of a previously uploaded tag, there are no visible changes to the system. However, when a user wants to see a tag cloud of the system, we have added some changes. When the user first enters our system, she is shown a tag cloud of the all-time-most-popular tags and the opportunity to narrow the time frame by choosing a year. If the user chooses a year, she is presented with the tags for that year and the possibility to further narrow the time frame by choosing a month. This narrowing of the time frame continues until the user has chosen a complete date (i.e. year, month and day). In addition to showing the user a tag cloud for each chosen time frame, we also show the user the time frame she is currently viewing. This we have done so the user will always be able to keep track of her chosen time frame. A picture of the interface is shown in figure 20.

In figure 21 we show the choices the user is presented with. As we can see, the user is only presented with the months and days that are actually available in the database, this is to prevent the user from specifying null sets e.g. choosing a month that results in an empty tag cloud because there are no tags in that month. The user is shown a new tag cloud each time she chooses a time frame (i.e. a year, month or day), and this is in following with the principle of tight coupling described in [60].

By showing the user a tag cloud for each chosen time frame, as well as before she has chosen a time frame at all, we have given the user the choice whether or not to manipulate the time frame. This way the graphical user interface (GUI) has not become much more complicated. There is however one thing that may complicate GUI; our system is very small and the lists of months and days are not very long. In a large CTS where tags and resources are added and changed every day, the list of months will contain 12

Figure 20: Our interface for choosing the time frame

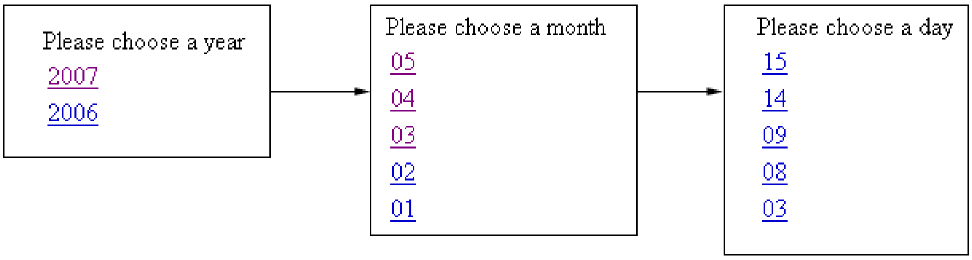

Figure 21: The graphical user interface that allows our user to narrow the time frame

choices and the list of days will contain 365 choices. With such large lists the GUI will be cluttered with information and the user may be overwhelmed. To avoid this, the dates can be shown as drop-down boxes instead of lists.

To summarize, we do not see the changes needed to implement our time-utilization as difficult to accomplish. The changes needed in the back-end of the system are mostly programming changes. And the changes on the front-end are merely design-issues as to how the user should choose the time frame. We do not believe that the additional changes to the graphical user interface, will complicate the user's tasks.

In the previous chapter, we evaluated time-utilization with the help of our pre-defined user roles. To summarize, we found that the familiar user was the one to benefit the most from time-utilization. With this functionality, she can monitor how the interests of the community change between each visit. And we stated that the unfamiliar user can employ a temporal tag cloud to gain a more thorough understanding of how the interests of the community shifts. By changing the time frame from which the tags are chosen, both the familiar and unfamiliar user can monitor if the community changes interests often or if it generally stays on one topic. With regards to our browsing and searching users, we found that their task neither benefit nor are hindered by the implementation of time-utilization. By implementing the functionality of time-utilization in our proof of concept, we have not discovered anything that changes our previous evaluation.

When comparing our proof of concept with a standard CTS, we do realize that there may be a difference in the results viewed by the user. In a large-scale CTS, the large amount of users and their tagging behaviour may be very complex. In this thesis, we will not be discussing how the interests of the users affect one another. The intention of our proof of concept is to show how time-utilization may be implemented. Because our system not is a collaborative system (the author has added the tags), we did not have the opportunity to analyse how our time-utilization enables a user to discover trends. The news archive of Gjøvik University College is not updated very often, and the articles posted are concerned with news about the school. The only trends we were able to discover was, that in periods, some topics have been discussed more often than others in periods (e.g. a competition that is mentioned in a couple of articles). We state that in a CTS, topics that continue over a longer time-period (e.g. the iraq war, or a presidential election), will affect the interests of the collaborating community — and thereby the tags used in the CTS. For future work, it could be interesting to see our proof of concept used on a collaborative system with a large body of users and resources.

5.3 Zooming functionality

When developing our proof of concept, we wanted to give the user the ability to obtain a more detailed view of the tag cloud. We implemented this with a zoom-functionality that allows the user to define the frequency-range required for a tag to be shown in the cloud. By setting the maximum required frequency for a tag, we give the user the ability to zoom in on the tag cloud. And by setting the minimum frequency, the user can limit the amount of words in the tag cloud. In figure 22 we show a flow chart of the zooming in our system.

In the previous section we described the part of our system that defines the time frame for our tag cloud. This part of the system takes the time frame as input and gathers the tags from the database. The system builds an array containing all the words in the

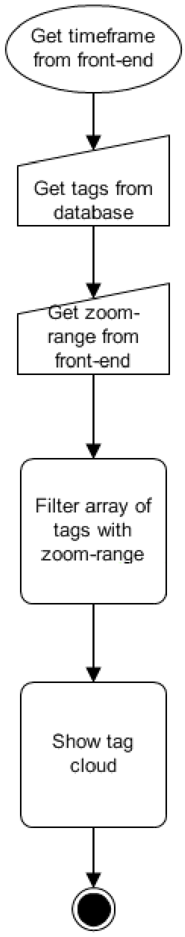

Figure 22: A flowchart of the zooming in our proof of concept

database (one entry for each word) and their frequency. If the user has entered a zoom-range, this array is then filtered according to this zoom-range.

When developing the user controls for the zooming, we considered letting the user drag a slider or to let the user enter the zoom range manually. We chose to let the user enter the minimum and maximum frequency manually. This choice is not following the principle of "Selection by pointing (not by typing)" from [60]. The reason we chose this method for setting the zoom range, is that it allows the user to enter both minimum and maximum frequency. To accomplish the same effect with sliders, the user would have to use two sliders and we decided that would be to many controls. In figure 23 we show the interface where the user enters the zoom-range.

5.3.1 Analysis

Now that we have shown how to implement zooming to enable the user to get a more detailed view, we will discuss the contributions of this implementation. By comparing our interface with that of a first generation tag cloud, we will look at the added functionality, as well as the increase in complexity.

Our system consists of two parts, as previously explained, the back- and the front-end. In the previous section we explained how the system starts by gathering the tags, and

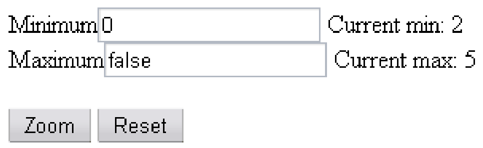

Figure 23: The interface to zoom in on our tag cloud

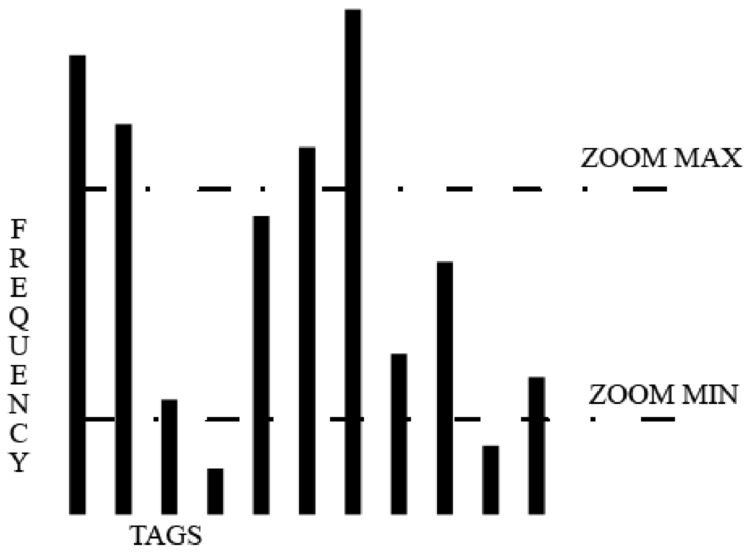

Figure 24: The zooming functionality of our proof of concept

their frequencies of use, from the database. The back-end then builds an array containing these data, and the system filters this array based on the zoom-range. For our small tagging system, this works very well. However, if the database in the tagging system becomes very large, like it will in a collaborative tagging system, this might put a lot of unnecessary strain on the system. For a large collaborative tagging system, we propose modifying the SQL-statements that gather the tags so that they only gather the tags within the zoom-range. This way the system will not have to traverse a large array; the array containing the tags will already be filtered.

In figure 23 we showed you the graphical user interface where the user enters the zoom-range. We also described how we had considered using a slider to set the zoom-range. We wanted to show that the zoom-range can be used to decrease the “height” of the tag cloud from the top, but also from the bottom. Figure 24 shows the principle of our zooming-functionality. Each bar in this histogram represents a tag, and its height represents its frequency. The upper-most dotted line is where the user has set the maximum frequency of the zoom-range, and the lower-most dotted line represents the minimum frequency of the zoom-range. Only the tags that have their frequency (i.e. if the bar stops) inside the zoom-range, will be shown.

By allowing the user to manipulate the required frequency for the tags in the cloud, we hope to enable her to gain more information from it. In our cloud the user can filter out the most popular tags, and gather information about which tags are, or are not that popular in the collaborative tagging system behind the tag cloud.

In [60], Ahlberg and Shneiderman describe the use of a starfield, or scatterplot, with zooming abilities. They state that this zooming can reduce cluttering. With our proof of concept, we hope to accomplish the same. We hope to enable the user to look past the largest and most domination tags in the cloud — thereby being able to get more information from the tag cloud.

On the front-end of our system, we have added two more controls to the interface (figure 23). With each new control we add to the interface, we run the risk of complicating the tag cloud-interface. To discuss the impact of our new controls, we firstly need to look at in which situations they will be used.

In the previous chapter, we discussed the implementation of zooming-functionality in the light of our four user-roles. These user-roles have been defined by their familiarity (unfamiliar and familiar) and the way they are looking for information (browsing and searching). We have previously explained how the unfamiliar user can employ the tag cloud to gain an impression of the underlying information. Because of the way the tag cloud emphasizes popularity, some topics are more visible than others. With our zooming-functionality, an unfamiliar user can gain more information not only about the most popular tags, but also about the topics that are not as popular. We state that the unfamiliar user can benefit from the zooming-functionality by getting a more in-depth knowledge of the tag cloud; She is not limited to seeing the "surface" and the largest tags.

When discussing the searching user, our proof of concept did not find that the zooming-functionality provided any specific benefits. When looking at the browsing user, we found that she may gain the ability to obtain more information from the tag cloud.

The browsing user has been stated to use the tag cloud to find words that she believes to be interesting. With the zooming-capabilities, it may be easier for her to find these words. On the "surface" of the tag cloud, the browsing user can easily see if there are any words that catches her eye. But if she does not find anything here, our zooming-functionality can enable her to look further into the tag cloud and continue her browsing there. These preliminary observations indicates that a tag cloud with zooming may be a more useful browsing tool than the first generation tag cloud.

As with the time-utilization, we would like to test our zooming-implementation on a large-scale collaborative tagging system (CTS). A large-scale test may provide us with more information on how a user can benefit from zooming. In a large-scale CTS, the tag cloud is a lot "deeper" (i.e. the frequency range of the tags is bigger), and the benefits from the in-depth view may be enhanced, as compared to our small-scale system.

We have previously explained how popularity can breed popularity and how the first generation tag cloud can be dominated by a few topics. By zooming in on the tag cloud, a user could ignore the dominating topics and the topics that are popular merely because of their own popularity. With this in mind, we believe that the tag cloud could provide a user with more interesting information, because the user is not limited to gain information about the most popular tags.

The previously described Tag Rolls from Delicious enables a user to manipulate the tag cloud of their own Delicious-tags, before publishing the tag cloud on their own home page. This is one situation where our functionality can be implemented; enabling a user to define the frequency-range of the tags in her cloud before publishing the finished tag cloud on her web page. In this situation we do not view our two new controls as a complication of the interface. We support this statement with the assessment that a user who is manipulating her tag cloud before publishing it, will not mind spending a little extra time to adjust it according to her liking.

The second situation where our zoom-controls can be implemented, is when a user is observing a tag cloud to gain information about the underlying information. In this situation, our new controls may be more complicating to the user. We do realize that the principle of the zoom-controls require some explanation. And one of the benefits of the tag cloud is its simplicity and ability to quickly and easily give the user an impression of the underlying information [10]. For a user to be able to operate our controls, she would have to be explained what they do. This added time does hinder the tag cloud as a quick visual information seeking tool. To be able to answer the question about whether or not our new functionality is practical, or if it just confuses the user, we will have to do user testing. This has not been done in this thesis and is left to future work.

As explained, we do realize that our zoom-controls may complicate the tag cloud-interface. To justify the implementation of them in the tag cloud-interface, we state that they apply new and needed functionality. In the introduction to this thesis, we stated that one of the limitations of the first generation tag cloud is that one or more topics tend to dominate the tag cloud. These dominating topics may hinder the user in finding a topic that interests her. With our proposed zoom-functionality, this user can zoom in on the tag cloud and ignore the largest tags. This way the dominating tags will be removed from the tag cloud enabling the user to more easily discover the smaller tags.

In the previous section, we described the Tag Rolls functionality. The Tag Rolls allows a user to adjust the size of her tag cloud and its fonts, as well as the font colours. In our proof of concept, we have not implemented this kind of functionality. The reason for this is that we do not consider this functionality to add any benefits with regards to the limitations that we stated in the introduction. The ability to manipulate the colour and size of the tag cloud, may be a nice functionality to adjust the visual appearance of the cloud. But we state that it does not improve the tag cloud as a visual information seeking interface.

5.4 Summary

In this chapter we have described the proof of concept that we have developed for this thesis. Our proof of concept implements time-utilization to enable the user to manually set the time frame for the tag cloud. Our research and analysis indicates that this time-utilization may be useful for the pre-defined "familiar" user-role. By comparing our system with a standard collaborative tagging system, we have done an assessment of the requirements for a system that implements time-utilization. We found that the database does not need any large changes, and neither does the system that gathers the tags for the cloud. The *relation-table* (figure 18) is necessary in both our system and a standard collaborative tagging system. The only difference is the introduction of the *time*-attribute that contains the date of the relation-creation, and this alteration should not increase the database to the extent where it becomes a problem.

With regards to the graphical user interface, our time-utilization demands some new controls. In our proof of concept, we have used hyperlinks to allow the user to choose a year, month or day. When the list of these time frames becomes very large, they will clutter the graphical user interface. To avoid this, we have proposed to use drop-down boxes instead. We have also discussed the use of a slider to dynamically change the time frame of the tag cloud.

Despite the issues of implementing time-utilization in the tag cloud, we state that it is a very useful functionality. It can help the user monitor which tags are popular in the collaborative tagging community, at a specific time frame.

Our proof of concept also implements zooming-functionality. We support the implementation of this on Ben Shneiderman's visual-information-seeking-mantra: "overview first, zoom and filter, then details on demand". We have compared our system with a standard collaborative tagging system, and shown that no difficult changes are needed. When filtering the tags for the tag cloud, we suggest two methods: Let the database collect all the tags and let the system filter them according to the zoom-range, or let the database filter the tags as it collects them. Our proof of concept is based on a very small database, and therefore we have not performed any testing on which of these methods are the most effective.

With our proof of concept, we hope to show how a second generation tag cloud may look and behave. We have shown some functionalities for a second generation tag cloud and how they may be implemented.

6 Summary and conclusion

Metadata is information about data. We use metadata to describe resources such as e.g. photos and web pages. By describing these resources with a set of descriptive words, we make it easier for ourselves, and others, to find them again later. In the last years, a type of method for assigning metadata has become popular — collaborative tagging. As opposed to the traditional method of assigning metadata where a professional assigns the metadata, collaborative tagging is about letting the users of the resources assign the metadata. This system of users, resources and tags, is called a collaborative tagging system.

As a collaborative tagging system evolves, it will eventually contain a very large body of information. An example of a popular collaborative tagging system is Delicious and, as of May 2007, they passed their two millionth user [73]. With such a large user base, the information on the collaborative tagging system may be difficult to browse through. To help a user find interesting resources, many systems employ an information visualization interface called a tag cloud. In this thesis we have been examining some limitations of this tag cloud, and ways of combating these limitations.

To help us evaluate the solutions we found to the limitations, and better understand how the tag cloud can be used as a visual information seeking interface,we defined a set of user-roles. These user-roles were described by their familiarity with the collaborative tagging system and how they are looking for information.

Unfamiliar user The unfamiliar user is a new visitor to a collaborative tagging system.

Familiar user The familiar user is a recurring visitor the collaborative tagging system. She has visited the system before, and has an understanding of the underlying information.

Browsing user The browsing user is browsing through the information in a system like a person might walk through a library looking at books to find something interesting. The user might have a general idea of the information she is looking for, but has not got a specific term or word for it.

Finding user The finding user has a specific term that she is looking for. Much like when a person searches on the Internet search engine Google

Some of the limitations we found were inherited from the underlying folksonomy. When the users of a collaborative tagging system work together, a common language will eventually form. This language is called a folksonomy, and this folksonomy is what the tag cloud visualizes. By examining the folksonomy, we found the origin of the limitations that are transferred to the first generation tag cloud. These limitations of the folksonomy and how they are transferred, are listed below.

Lack of structure Because of the lack of structure in a folksonomy, there is also no structure in the first generation tag cloud. In a folksonomy, as opposed to a taxonomy, there is no defined relationship between the tags.

Dominating topics Because there is no higher authority to maintain the tags in a folksonomy, there is a certain amount of inconsistency. This inconsistency manifests itself in the form of synonyms. When a topic gains popularity in the collaborative tagging system, this topic and all its synonyms can end up dominating the tag cloud. This limits the visibility of other interesting topics.

By studying research on folksonomies and tag clouds, we found ways to combat these limitations. Despite the lack of structure in the folksonomy, we found research that has been able to cluster tags according to their similarity. This can be done either with co-occurrence analysis or with latent semantic analysis. We showed a clustered tag cloud (figure 9), and how it had managed to cluster similar tags. This clustering limited the synonym-problem and visualized a structure between the tags in the cloud. In addition to clustering tags, we also described research that has managed to infer a hierarchy on the tags in a folksonomy.

By relating these results to our user-roles, and the two first limitations of the first generation tag cloud, we found indications that the both the unfamiliar and the familiar browser can benefit from clustering. We base this conclusion on the point that it is easier to understand the meaning of the tags in the cloud, and that the clustered tag cloud no longer is as dominated by a few topics. Our conclusion is also based on the statement that browsing benefits from an underlying structure [25].

The other limitations, of the first generation tag cloud, that we found were caused by the design of the tag cloud itself. We state that the way the first generation tag cloud is designed, it is provided with some limitations that inhibits it as an information visualization interface. These limitations are listed below.

Lack of time control The static design of the first generation tag cloud does not visualize how the underlying information changes over time. The underlying folksonomy directly relays the interests of the tagging community, and to monitor these changes can be interesting. By monitoring these changes, we can discover trends as well as see how tags increase or decrease in popularity.

Lack of a detailed view of the tags The most popular tags in a cloud is very visible to the user, but she does not have the possibility to get a more detailed view of the less popular tags.

Very wide search-field When a user finds a word in the tag cloud that she finds interesting, she clicks on this word. As she clicks on the word, the user is presented with all the resources that have been tagged with this word. As a search query, one word is a very wide search, and the result is a very large result set. In this result set, there will be many resources that are not related to the what the user was looking for.

Lack of interactivity In the first generation tag cloud there are not many ways a user can manipulate the presentation of the tag cloud. The user viewing the tag cloud has to accept the organization of the tags in the cloud, the font sizes, colours and so forth.

By studying related literature, we explained how we can benefit from the ability to monitor the temporal changes in a collaborative tagging system. We described how the monitoring of temporal changes can help us gather information on how the interests of the community shifts over time. The only system we found that utilized time the way we visualized it, was Tag Lines, so to further examine the effects of this implementation, we developed a proof of concept. By relating this literature, and our proof of concept, to our user-roles, we found indications that the familiar user can benefit from the utilization of time in the tag cloud. With a tag cloud that implements time-utilization, the familiar user can see how the information in the collaborative tagging system has changed since her last visit.

When stating that the lack of a detailed view of the tag cloud was a limitation, we based this on Shneiderman's visual-information-seeking-mantra; "overview first, zoom and filter, then details on demand". We were not able to find any implementations that allowed the user to zoom and filter on the tag cloud, so we implemented a zooming-functionality in our proof of concept. With the help of our proof of concept, we state that our user-roles can benefit from the ability to zoom past the largest tags and gain a more detailed view of the smaller tags. We state that our zooming-functionality to some extent can help combat the problem of dominating topics; as the user can zoom past them.

When studying the related literature, we found a way to avoid the wide one-word-search-field of the first generation tag cloud. We found that faceted browsing can help our finding user in narrowing and specifying her search-query. With this in mind, we conclude that faceted browsing improves the tag cloud with regards to finding information.

We have stated that the lack of interactivity in the first generation tag cloud is mainly because there are not any alternatives to the standard presentation. However, the related literature showed us that the presentation of the first generation tag cloud can affect the usability and effectiveness of it. Our conclusion is that as the second generation tag cloud implements more functionality, the user could benefit from the ability to choose how it is presented to her. An example of this is the choice as to whether or not the tags in the cloud should be arranged alphabetically or with clusters.

When we implemented time-utilization and zooming-functionality in our proof of concept, we wanted to evaluate their impact on both the system behind the tag cloud and the tag cloud-interface. We did not perform any user studies on our proof of concept. Instead we discussed how the implementation of more user-controls may affect the tag cloud as an easy-to-use interface. As stated, one of the benefits of the tag cloud is its simplicity, and we did not want to compromise this.

With regards to the implementation of time-utilization, we found that there are some new requirements to the underlying system, but nothing very problematic. In the graphical user interface, we did find some problems with our time-utilization-controls. To avoid these, we proposed other ways of allowing the user to set the time frame for the tag cloud.

When we analysed the zooming-functionality of our proof of concept, we found that it could easily be implemented in a standard collaborative tagging system. The system for choosing tags and building the tag cloud did not require many alterations to enable our zooming. When analysing the impact, of the zooming-controls on the graphical user interface, we found that the controls can be implemented without many problems. We did however realize that the principle of the zooming can be difficult and time-consuming to comprehend for the user. To avoid this, we proposed to implement the controls as a slider — much like the zooming-control on an electronic map.

Our overall conclusion is that a second generation tag cloud, with more functionality, can help restrict the limitations of the first generation tag cloud. We state the overall contribution of the functionality shown in this thesis can improve the tag cloud as a browsing tool, and that it can help both a familiar and unfamiliar user gain more information from the visualization. Our conclusions are not based on user studies, but on analysis of the limitations of the first generation tag cloud and the contributions of the suggested functionalities. With this thesis, we hope that we have provided an insight into the visualization of collaborative tagging information with the help of tag clouds. In addition to evaluating the tag cloud, we also wanted to “create an order in the chaos” surrounding tag clouds and their relation to folksonomies and browsing.

7 Future work and implementations

7.1 Future implementations

We have described how our proof of concept enables the user to narrow the time frame of the tag cloud by choosing year, month and day. We have also described how the graphical user interface can be cluttered with links in our visualization. To avoid this we propose showing the years, months and days in drop-down boxes.

With our implementation of time-utilization, the user is restricted to choosing one year, month or day for the time frame. A different interface could be developed to allow the user to choose multiple years, months or days.

In addition to our implementation of time-utilization, we would like to include a slider to allow the user to change the time frame and dynamically see the tag cloud change. We propose a slider with a set start-date and end-date. By dragging the slider the user can change either the start- or end-date and thereby narrowing the time frame. As the slider is dragged, some tags grow in size as others diminish and disappear, and the user can dynamically monitor the changes. To do this the system will have to implement dynamic communication between the front- and back-end, one way of doing this is with Asynchronous JavaScript and XML (AJAX) [74]. We tried to accomplish this, but we had programming problems when trying to pass multiple variables between the front- and the back-end of the system. Unfortunately, the time available for this thesis did not allow us to fix this problem. So we left that functionality out of the proof of concept.

In the previous section, we mentioned the problem of increasing the complicity of the tag cloud-interface. To simplify the zoom-controls, we propose a different way of implementing the zoom-functionality. As we show in figure 24, our zoom-functionality manipulates both the minimum and maximum frequency of the tags in the cloud. A different implementation could be to apply a fixed distance between the minimum and maximum frequency. This way the user would only need one control to move the maximum frequency. This way we can halve the new number of controls, and we believe that the functionality will be a lot easier to explain. This control could be implemented as a slider at the side of the tag cloud. This slider can easily be compared to the zoom-control of e.g. Google Maps [75]. We state that this zoom-control (figure 25), is easily recognized by a user who has used an electronic map before.

Previously we described faceted browsing and how it can help the user narrow her search-field. Unfortunately the time available for this thesis, did not allow us to implement this. For future implementations, we propose that when a user clicks on a tag, she is shown both the related resources and a new tag cloud. This new tag cloud will only contain tags that have been used in conjuncture with the previously chosen tag. We propose that this process is repeated until the amount of tags in the cloud recede below a certain number. This way the user will be presented with the related resources at every step, in addition to the ability to refine her search. We propose that faceted browsing, can aid the user-role we previously defined as a “searching user”, in finding the information she is looking for.

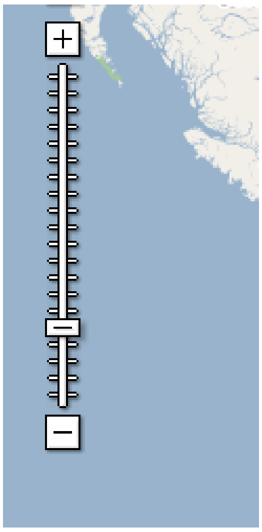

Figure 25: The zooming-controls from Google Maps

Unfortunately, we were not able to develop a tag cloud with all our envisioned functionalities. To compensate we have a vision for the graphical user interface of a second generation tag cloud. This is shown in figure 26. This figure shows the graphical user interface of what we envision as a complete second generation tag cloud.

As we can see, it implements a slider to dynamically change the time frame of the tag cloud. What we wanted to develop was an interface where the user can set the time frame with the years or months listed, and then use the slider to navigate within that time frame.

To the left of the tag cloud, we can see the easily recognizable zoom-slider from an electronic map. This slider enables the user to zoom in on the tag cloud. It functions as we have described above; with a set range between the maximum and minimum frequency of the tags in the cloud.

On the right side of the interface, we can see a check box. With this check box, the user can choose if she wants her tags arranged alphabetically or clustered. To the right of the interface, we can show a column labelled "Related resources". This column shows resources that have been tagged with the word the user clicks in the cloud. The tag cloud itself changes to show words that have been used in conjuncture with the previously chosen tag.

The interface may look small and cramped in our figure, but on a computer screen it will be bigger. We state that a 1024x768 pixels is more than enough for visualizing this interface without the controls being cramped to much together. According to W3Schools, 80% of their users have a screen resolution of 1024x768 or higher [76] (as of January 2007). The statistics of this site are not completely representative, as the users are generally interested in web technologies. In web design, it is generally a rule of thumb to design a web page that is viewable at all screen resolutions — without the user having to scroll horizontally. At a screen resolution of 800x600 pixels, the interface will use up a lot of the available room on the screen. It should however not be necesarry for the user to scroll to view the entire interface. This does depend on the size of the tag cloud, but that can be adjusted by the designers of the system.

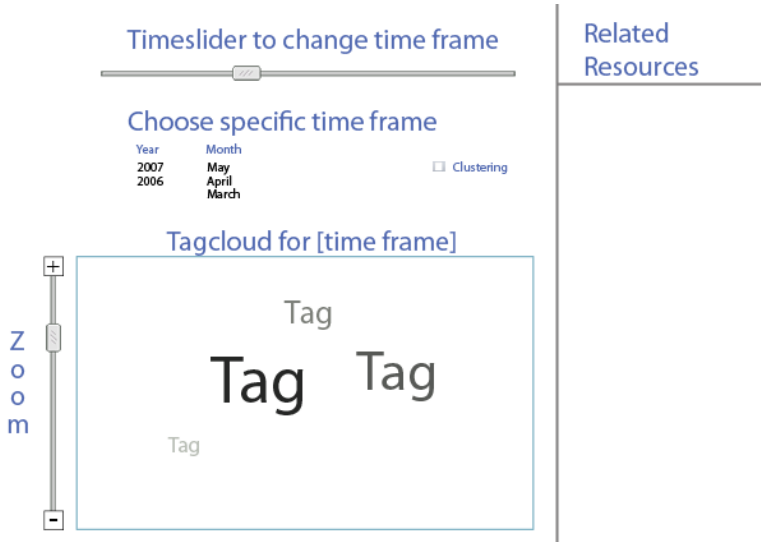

Figure 26: Our vision of a second generation tag cloud

Figure 26 shows how we wanted our prototype of a second generation tag cloud. As we can see the interface has a great deal more controls. This addition of controls reduces the simplicity of the first generation tag cloud, and that is undesirable. We state that if the tag cloud is to increase in popularity and usability, it needs these new controls and functionalities to allow the users to gain more information from the visualization. However, we can not fully ascertain whether or not these controls inhibits the user or makes the process of information seeking easier, without doing user studies.

7.2 Future work

In the previous section, we described our intentions for future implementations in a second generation tag cloud. This section will describe future work on tag clouds in general.

As a suggestion to altering tag clouds, Montero and Solana proposed to cluster the tags in the cloud [9]. We would like to see user studies performed to evaluate the effectiveness of a clustered tag cloud. Halvey and Keane performed some user testing in [21], but they looked at the position and size of the tags in the cloud. In addition, they gave their participants goal-oriented tasks as opposed to free-browsing tasks. We propose that the tag cloud is mainly a browsing interface and that user studies on the tag cloud should be done with free-browsing tasks.

As described in our thesis, there are several suggestions as to how the clustering can be performed. We would like to see some analysis as to how effective these clustering techniques are. By effective, we refer both to the clusters they produce, and their ability to cluster the tags in real-time — so that the tag cloud can continue as a simple and dynamic interface. And, as mentioned by Halvey and Keane in [21], there has not been any evaluation as to whether the users find any improvements in the clustered tag cloud.

We previously mentioned how Graham et al states the importance of ensuring "real *user involvement* in the design process". In accordance with this, we state that we need to perform more user studies on the tag cloud. These user studies should focus on how the tag cloud performs as a browsing interface. And the users should be tested on tag clouds both with and without new functionalities to evaluate if the new functionalities are hindering or helping the users in completing their assigned task. We propose to perform user studies with a focus on browsing task, as we have concluded that the tag cloud is good tool for browsing. We would also like to see studies performed on the user-roles we have employed for the evaluations in this thesis. With the help of our user-roles, we believe that researchers can easily define user-situations and tasks to use in future user studies.

Bibliography

[1] Flickr — photo sharing. http://www.flickr.com/. Online; accessed June 30 - 2007.

[2] del.icio.us. http://del.icio.us/. Online; accessed June 30 - 2007.

[3] Technorati:home. http://www.technorati.com/. Online; accessed June 30 - 2007.

[4] Golder, S. & Huberman, B. A. April 2006. The structure of collaborative tagging systems. *Journal of Information Science*, 32(2), 198–208.

[5] Sinha, R. September 2005. A cognitive analysis of tagging. http://www.rashmisinha.com/archives/05_09/tagging-cognitive.html. Online; accessed June 30 - 2007.

[6] Wu, X., Zhang, L., & Yu, Y. 2006. Exploring social annotations for the semantic web. In *WWW '06: Proceedings of the 15th international conference on World Wide Web*, 417–426, New York, NY, USA. ACM Press.

[7] Mathes, A. Folksonomies - cooperative classification and communication through shared metadata. Online; accessed June 30 - 2007, December 2004.

[8] Google. http://www.google.com. Online; accessed June 30 - 2007.

[9] Hassan-Montero, Y. & Herrero-Solana, V. 2006. Improving tag-clouds as visual information retrieval interfaces. In *Proceedings of the International Conference on Multidisciplinary Information Sciences & Technologies*.

[10] Rivadeneira, A. W., Gruen, D. M., Muller, M. J., & Millen, D. R. 2007. Getting our head in the clouds: toward evaluation studies of tagclouds. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, 995–998, New York, NY, USA. ACM Press.

[11] Tagcrowd. http://www.tagcrowd.com/. Online; accessed June 30 - 2007.

[12] Pattern recognition » tag cloud for 2007 state of the union. http://www.jasongriffey.net/wp/2007/01/24/tag-cloud-for-2007-state-of-the-union/. Online; accessed June 30 - 2007.

[13] Begelman, G., Keller, P., & Smadja, F. 2006. Automated tag clustering: Improving search and exploration in the tag space. In *WWW2006: Proceedings of the Collaborative Web Tagging Workshop*, Edinburgh, Scotland.

[14] Porter, J. October 2005. Web2con: Popularity decay in tagging. http://bokardo.com/archives/web2con-popularity-decay-in-tagging/. Online; accessed June 30 - 2007.

[15] Graham, M., Kennedy, J., & Benyon, D. 2000. Towards a methodology for developing visualizations. *Int. J. Hum.-Comput. Stud.*, 53(5), 789–807.

[16] Acm digital libray. http://www.acm.org/dl/. Online; accessed June 30 - 2007.

[17] Ieee xplore: Dynamic home page. http://www.ieeexplore.ieee.org/. Online; accessed June 30 - 2007.

[18] Gershon, N., Card, S., & Eick, S. G. 1999. Information visualization tutorial. In *CHI '99: CHI '99 extended abstracts on Human factors in computing systems*, 149–150, New York, NY, USA. ACM Press.

[19] Shneiderman, B. & Plaisant, C. Mar 2004. *Designing the User Interface: Strategies for Effective Human-Computer Interaction*. Addison Wesley, 4th edition edition.

[20] Skaalid, B. Human-computer interface design. http://www.usask.ca/education/coursework/skaalid/theory/interface.htm. Online; accessed June 30 - 2007.

[21] Halvey, M. J. & Keane, M. T. 2007. An assessment of tag presentation techniques. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, 1313–1314, New York, NY, USA. ACM Press.

[22] Plaisant, C. 2004. The challenge of information visualization evaluation. In *AVI '04: Proceedings of the working conference on Advanced visual interfaces*, 109–116, New York, NY, USA. ACM Press.

[23] Kuo, B. Y.-L., Hentrich, T., Good, B. M. ., & Wilkinson, M. D. 2007. Tag clouds for summarizing web search results. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, 1203–1204, New York, NY, USA. ACM Press.

[24] Morse, E., Lewis, M., & Olsen, K. A. 2000. Evaluating visualizations: using a taxonomic guide. *Int. J. Hum.-Comput. Stud.*, 53(5), 637–662.

[25] Lin, X. 1997. Map displays for information retrieval. *J. Am. Soc. Inf. Sci.*, 48(1), 40–54.

[26] Halpin, H., Robu, V., & Shepherd, H. 2007. The complex dynamics of collaborative tagging. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, 211–220, New York, NY, USA. ACM Press.

[27] Power law — wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Power_law. Online; accessed June 30 - 2007.

[28] Mesnage, C. & Jazayeri, M. 2006. Specifying the collaborative tagging system. In *SAAW'06, Semantic Authoring and Annotation Workshop*.

[29] Furnas, G. W., Landauer, T. K., Gomez, L. M., & Dumais, S. T. 1987. The vocabulary problem in human-system communication. *Commun. ACM*, 30(11), 964–971.

[30] The esp game: Labeling the web. http://www.espgame.org. Online; accessed June 30 - 2007.

[31] von Ahn, L. & Dabbish, L. 2004. Labeling images with a computer game. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, 319–326, New York, NY, USA. ACM Press.

[32] Quintarelli, E. June 2005. Folksonomies: Power to the people. http://www.iskoi.org/doc/folksonomies.htm. Online; accessed June 30 - 2007.

[33] Sen, S., Lam, S. K., Rashid, A. M., Cosley, D., Frankowski, D., Osterhouse, J., Harper, F. M., & Riedl, J. 2006. tagging, communities, vocabulary, evolution. In *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, 181–190, New York, NY, USA. ACM Press.

[34] Li, R., Bao, S., Yu, Y., Fei, B., & Su, Z. 2007. Towards effective browsing of large scale social annotations. In *WWW '07: Proceedings of the 16th international conference on World Wide Web*, 943–952, New York, NY, USA. ACM Press.

[35] Lawrence block — wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Lawrence_Block. Online; accessed June 30 - 2007.

[36] Desire lines — wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Desire_lines. Online; accessed June 30 - 2007.

[37] Merholz, P. October 2004. Metadata for the masses. Online; accessed June 30 - 2007.

[38] Marlow, C., Naaman, M., Boyd, D., & Davis, M. 2006. Ht06, tagging paper, taxonomy, flickr, academic article, to read. In *HYPERTEXT '06: Proceedings of the seventeenth conference on Hypertext and hypermedia*, 31–40, New York, NY, USA. ACM Press.

[39] Cattuto, C., Loreto, V., & Pietronero, L. January 2007. From the cover: Semiotic dynamics and collaborative tagging. *PNAS*, 104(5), 1461–1464.

[40] Cattuto, C., Loreto, V., & Pietronero, L. May 2006. Collaborative tagging and semiotic dynamics.

[41] Catutto, C., Schmitz, C., Baldassarri, A., Servedio, V. D. P., Loreto, V., , Hotho, A., Grahl, M., & Stumme, G. 2007. Network properties of folksonomies. *AI Communications Journal, Special Issue on "Network Analysis in Natural Sciences and Engineering"*.

[42] Michlmayr, E. & Cayzer, S. 2007. Learning user profiles from tagging data and leveraging them for personal(ized) information access. In *Proceedings of the Workshop on Tagging and Metadata for Social Information Organization, 16th International World Wide Web Conference (WWW2007)*.

[43] Mika, P. 2007. Ontologies are us: A unified model of social networks and semantics. *Web Semant.*, 5(1), 5–15.

[44] Choy, S.-O. & Lui, A. K. 2006. Web information retrieval in collaborative tagging systems. *wi*, 0, 352–355.

[45] Landauer, T. K., Foltz, P. W., & Laham, D. 1998. An introduction to latent semantic analysis. *Discourse Processes*, 25, 259–284.

[46] Gorrell, G. May 2005. Latent semantic analysis: How does it work, and what is it good for? http://www.dcs.shef.ac.uk/˜genevieve/lsa_tutorial.htm. Online; accessed June 30 - 2007.

[47] Brooks, C. H. & Montanez, N. 2006. Improved annotation of the blogosphere via autotagging and hierarchical clustering. In *WWW '06: Proceedings of the 15th international conference on World Wide Web*, 625–632, New York, NY, USA. ACM Press.

[48] Technorati: About us. http://www.technorati.com/about/. Online; accessed June 30 - 2007.

[49] Porter, J. October 2005. Web2con: Emergent tags. http://bokardo.com/archives/web2con-emergent-tags/. Online; accessed June 30 - 2007.

[50] Dubinko, M., Kumar, R., Magnani, J., Novak, J., Raghavan, P., & Tomkins, A. 2006. Visualizing tags over time. In *WWW '06: Proceedings of the 15th international conference on World Wide Web*, 193–202, New York, NY, USA. ACM Press.

[51] Taglines. http://research.yahoo.com/taglines/. Online; accessed June 30 - 2007.

[52] Cloudalicious — watching tag clouds over time. http://cloudalicio.us/. Online; accessed June 30 - 2007.

[53] Russell, T. 2006. cloudalicious: folksonomy over time. In *JCDL '06: Proceedings of the 6th ACM/IEEE-CS joint conference on Digital libraries*, 364–364, New York, NY, USA. ACM Press.

[54] Rech, J. 2007. Discovering trends in software engineering with google trend. *SIGSOFT Softw. Eng. Notes*, 32(2), 1–2.

[55] Nowell, L. T., France, R. K., Hix, D., Heath, L. S., & Fox, E. A. 1996. Visualizing search results: some alternatives to query-document similarity. In *SIGIR '96: Proceedings of the 19th annual international ACM SIGIR conference on Research and development in information retrieval*, 67–75, New York, NY, USA. ACM Press.

[56] Google trends. http://www.google.com/trends. Online; accessed June 30 - 2007.

[57] Popular tags on flickr photo sharing. http://www.flickr.com/tags/. Online; accessed June 30 - 2007.

[58] del.icio.us tagrolls. http://del.icio.us/help/tagrolls/. Must be logged in to view - Online; accessed May 30 - 2007.

[59] Zoomclouds. http://zoomclouds.egrupos.net/. Online; accessed June 30 - 2007.

[60] Ahlberg, C. & Shneiderman, B. 1994. Visual information seeking: tight coupling of dynamic query filters with starfield displays. In *CHI '94: Proceedings of the SIGCHI conference on Human factors in computing systems*, 313–317, New York, NY, USA. ACM Press.

[61] Kaser, O. & Lemire, D. May 2007. Tag-cloud drawing: Algorithms for cloud visualization. In *WWW '07: 16th International World Wide Web Conference*. To appear in proceedings of Tagging and Metadata for Social Information Organization (WWW 2007).

[62] Bielenberg, K. & Zacher, M. Groups in social software: Utilizing tagging to integrate individual contexts for social navigation. Master's thesis, Program of Digital Media, Universität Bremen, 2005.

[63] Shaw, B. December 2005. Utilizing folksonomy: Similarity metadata from the del.icio.us system. http://www.metablake.com/webfolk/web-project.pdf. Online; accessed June 30 - 2007.

[64] Westerman, S. J. & Cribbin, T. 2000. Mapping semantic information in virtual space: dimensions, variance and individual differences. *Int. J. Hum.-Comput. Stud.*, 53(5), 765–787.

[65] Sutcliffe, A. G. & Patel, U. 1996. 3d or not 3d: Is it nobler in the mind? In *HCI '96: Proceedings of HCI on People and Computers XI*, 79–94, London, UK. Springer-Verlag.

[66] Sinha, R. July 2006. Findability with tags: Facets, clusters, and pivot browsing. http://www.rashmisinha.com/archives/06_07/tag%20findability.html. Online; accessed June 30 - 2007.

[67] Yee, K.-P., Swearingen, K., Li, K., & Hearst, M. 2003. Faceted metadata for image search and browsing. In *CHI '03: Proceedings of the SIGCHI conference on Human factors in computing systems*, 401–408, New York, NY, USA. ACM Press.

[68] Wine.com — buy wine, wine clubs, gift baskets and more. http://www.wine.com/. Online; accessed June 30 - 2007.

[69] Rawsugar — home page. http://www.rawsugar.com/. Online; accessed June 30 - 2007.

[70] Mysql ab::a guide to database high availability. http://www.mysql.com/. Online; accessed June 30 - 2007.

[71] English — gjøvik university college. http://english.hig.no/. Online; accessed June 30 - 2007.

[72] Php: Hypertext preprocessor. http://www.php.net/. Online; accessed June 30 - 2007.

[73] del.icio.us: that was fast. http://blog.del.icio.us/blog/2007/03/that_was_fast.html. Online; accessed June 30 - 2007.

[74] Ajax (programming) — wikipedia, the free encyclopedia. http://en.wikipedia.org/wiki/Ajax_(programming). Online; accessed June 30 - 2007.

[75] Google maps. http://maps.google.com/. Online; accessed June 30 - 2007.

[76] W3schools - browser display statistics. http://www.w3schools.com/browsers/browsers_display.asp. Online; accessed June 30 - 2007.

A PHP-code and database-definitions

If the reader wishes to see the PHP-code used in this thesis, it can be downloaded as a rar-archive at `http://www.mgns.no/theses/`. Below follows a description of the files in the rar-archive.

index.php This is the file that builds the graphical user interface. It includes three class-definition-files (described below).

style.css Contains style-definitions for the graphical user interface.

database.sql Contain SQL for the definition of the database and its contents. The file has been created with the help of phpMyAdmin.

classdef/databaseconnection.php Contains the class-definition for the class *databaseconnection*, that handles communication to the database.

classdef/gui.php Contains the class-definition for the class *gui*, that handles the selection and collection of tags.

classdef/wordcloud.php Contains the class-definition for the class *wordcloud*, that builds and filters the tag-array and prints the HTML-code to visualize the tag cloud. This file is based on a script by Derek Harvey (`www.derekharvey.co.uk`).

The rar-archive can be unpacked with, among others, WinRAR (a trial version can be obtained from `http://www.rarlab.com/download.htm`).