UNIVERSITY OF OSLO
Department of Informatics

Collaborative Symbolic Structuring in a Polyscopic Environment

Master thesis (60 credits)

Morten Warncke-Wang

August 1, 2008

Abstract

The problem of information overload was described by Vannevar Bush as early as 1945 in his article “As We May Think”, and as we continue to make it easier to create information we also increase the probability of getting lost in the maze. We seek to improve the quality of the information we create and the way we access it by combining the associative paths of Vannevar Bush, the symbolic structures of Douglas Engelbart and the principles of Polyscopy by Dino Karabeg, and show how these can all come together to form an untraditional approach to organizing information on the Web.

Those who cannot remember the past
are condemned to repeat it.
(George Santayana, The Life of Reason, 1905)

We all benefit from experience;
preferably someone else's
(Footprints project slogan)

Contents

1 The Problem

Our research has discovered not only that information overload is a problem, but that it's also one that has been talked about for more than 60 years. It comes as no surprise that we're having problems coping with too much information today when for instance there's approximately ten *hours* of video uploaded to YouTube every *minute*[42], but if this was a problem half a century ago, and one that hasn't really improved since then, what have we been doing in the meantime?

Although computers are a fairly recent invention, information technology is an old discipline. Many of the problems we are working on have been pondered on for centuries as they're not fundamentally computer-related. For instance classification systems and indexing are the work of librarians, a problem they have been faced with since the days of scrolls. In today's society where the Internet is everywhere the search engine is a click or two away and allows us instant retrieval. Where we earlier had to go back home an consult our encyclopedias to resolve arguments we now pick up our cell phones and google it. With it also comes access to the latest news from anywhere around the globe, updated only minutes after it happens. Actually it's not just anywhere around the globe, it's anywhere in our solar system, the latest images of water on Mars are also available on your cell phone.

We don't see this as a bad thing, having a world of information available everywhere. The Web allows both quick publishing and retrieval, but do we always find what we're looking for? A search engine might give you a million web pages matching your query, but how many of them did you read and find useful? Or did you follow the first match, notice a few other interesting links there, and end up spending an hour reading comics instead of what you were supposed to be doing? Can we, to borrow a metaphor from the movie Spinal Tap where the guitarist had an amplifier where the controls go all the way to 11, make a solution that goes “one louder”?

1.1 The Idea

This thesis tarted out as an idea of attaching a semantic network on top of a Wiki. We wanted a human approach to creating that semantic network, in many ways similar to how Wikipedia works. Not because collaborative websites are in fashion right now, but instead because we believe people coming together to work on something can make it better, utilize the Wisdom of the Crowd as James Surowiecki calls it[41]. Through cooperation complex things as an operating system, an office suite and an encyclopedia has been created. We're referring to Linux, OpenOffice.org and Wikipedia, respectively. Although open source software and volunteer work has a business model that's a lot unlike what we're used to, they are all success stories showing that it can work given the right circumstances.

Another interesting example is the bottom-up approach to classification that takes place through what is called "tags". Keywords that you can use to describe a photo you uploaded to Flickr, a bookmark you added to del.icio.us, or a blog post you notified Technorati about. The vocabulary is unrestricted, it allows anyone to describe what they have using the words that fit for them, which means tags can be used a mnemonic devices, they aid you in retrieving what you stored[6]. You would've thought that the freedom in this kind of systems makes it useless as nobody uses common language to describe items, but actually the vocabulary converges over time as people use the system, partly because the descriptions used by others affect the words you choose[39, 38]. The beauty still remains though, that you're free to choose your own description.

Can we combine this bottom-up approach to create a semantic network and integrate it into a Wiki? Part of the answer must be yes as research reveals that semantic wikis already exists. We could've gone the same route, learned from their experiences and attempt to move the technology a step forward. Would that result in a system that goes "all the way to 11"?

Rather than integrate the semantic network with the Wiki technology, we found it a lot more interesting to see if it could be decoupled from it. Making a system that could be hooked up to just about any website sounded a lot more interesting. And could we come up with a system that enables users to manipulate a graphic representation of the network, rather than go with textual representation? This untraditional approach to how we organize information on the Web sounded tempting. It was also an idea that we found would realize the visions of famous scientists like Vannevar Bush and Douglas Engelbart. It meant we had to give it a go.

1.2 The Structure

This thesis does not follow the traditional structure, instead we have chosen to split it up into three main parts. These parts are interconnected, they come together to create a complete structure. First is the foundation, background that is the basis for the parts that follow. Upon the foundation we will create a rough pencil sketch from connecting some bigger dots and see some connections between the pieces in the foundation that might be missed on first pass. Lastly we will bring an eraser to the sketch and replace it with clear drawings of the ideas on how to realize our vision.

Througout this thesis you will see references to pieces of information that we have found interesting and relevant, you might already have noticed a few of them. Our work on this research has provided us with diverse resources, there is scientific articles, popular science books, and blogs and websites. Rather than limit the amount of references we have chosen to use them all so that you might check them out yourself and find something interesting too. Additionally you will find that the knowledge we review as part of our foundation are what we regard as the really interesting bits.

At the end of the thesis we will round it all up and look at alternatives for where to go next. Through the work with this thesis we have discovered ideas that have potential to be put to the test. We started by mentioning that information overload has been a problem for more than 60 years, and as we will see near the end, some really cool ideas in information technology comes from the beginning of the 20th century.

2 The Foundation

We start out with the foundation which we have chosen to do thematically. The first theme is organization of information which we will cover by exploring three articles. The second theme is the main technology for the solution, the World Wide Web, which we will cover both from a historical perspective, show some important learned lessons to bring with us, as well as show that the basis on which the Web was built might not be as shallow as you might think.

2.1 As We May Think

The first piece of my foundation is a classic article, “As We May Think”[4] by Vannevar Bush, published in the Atlantic Monthly in 1945. He is probably most known for this article, but many of the ideas he put forth can also be found in the essay “Mechanization and the Record” from 1939. In between the publication of the two articles were the Second World War, and the impact of the war can be seen. It is the end of the war when Bush writes “As We May Think” and he is asking what ideas will be occupying the minds of engineers, scientists and thinkers for the years to come. Not just any kind of ideas, he points out that the question is now what projects will take them to the edge of their abilities, keep them up at night and bring out the best in themselves?

> “Now, as peace approaches, one asks where they will find objectives worthy of their best.”

During the years the war had lasted, engineers had come up with some ingenious solutions to help fight the Third Reich. Two examples are the bouncing bomb used by the 617th Squadron to attack German dams on May 17 1943 in Operation Chastise, by having the bomb skip on the water to reach the dam rather than being caught in the mine nets. Another example is the Mulberry harbours built by the allied forces off the coast of Normandy on D-Day in 1944, which not only had to be created in such a way that they could be shipped in from England, but they also had to work with tidal water changes of 21-24 feet found at those sites. Many more examples can be found and it should come as no surprise that Vannevar Bush is asking what future projects will motivate similar perservering effort and creativity.

2.1.1 Information Magnitude

Remember that the calendar has not yet reached half way through the twentieth century when "As We May Think" was published, but already at that time the amount of information available was an issue. Maybe it was more an issue of available information technology as it attempted to keep up with demand. Either way the amount of information was already so large that it forced scientists to specialize, and as Bush argues that specialization in turn allows us to create more information as we dig deep towards the core of our field. This means we're looking at a problematic spiral that results in an ever increasing pile of information which puts a strain on information technology because, as Bush points out, if information is to be really useful for us it has be a dynamic entity:

> "A record if it is to be useful to science, must be continuously extended, it must be stored, and above all it must be consulted."

This concept is maybe so simple that it is obvious, and it leaves us with three challenges. First of all it is pointing out that we need to be able to retrieve that which is already stored, our information systems must make us able to find the information we are looking for, and by that we mean not just finding what we might think we are looking for. Secondly we must be able to extend that which already exists rather than having to replicate it, which is why it is important to find what we're looking for. Lastly, when we're done we must be able to store it in a way that enables everyone to find it, extend it, and store it again. It is a continuous cycle of adding various amounts of drops to our pool of knowledge. Bush mentions the example of Gregor Mendel, a monk whose concepts of the law of genetics were disregarded for many years as he was not a reputable scientist, an example that we must take care to avoid repeating. Can we create systems that make us able to find and utilize that which we are seeking, although we might be uncertain of what it is we're searching, as well as bringing forward the voice of those with good ideas instead of drowning them in noise from a crowd of shouters?

2.1.2 More than one way to Rome

In his article, Bush describes the job of navigating through information as something that has not improved much over the past century, he compares it to sailing on the high seas with square-rigged ships. The field of information storage and retrieval has of course made huge improvements since 1945, with today's database capabilities we can now search through millions of records in next to no time. We are also getting better at harnessing human knowledge as we described in the introduction to decide what is relevant to us, something which we will explore in more detail in section 3.2.1.

The issue mentioned in the quote in section 2.1.1 is important to Bush partly because at the time of writing the ability to retrieve information was limited. To find something at the library you might have to thumb through an index sorted alphabetically by title, or by author, or navigating by a classification system. The shortest route could still be the human one, instead of looking it up yourself you'd ask the librarian, just like people still ask each other rather than searching for something. The common thing here is that there needs to be many ways to reach the same result, finding the information you're looking for, there needs to be more than one way to Rome. In the book "About Face"[5] Alan Cooper sums it up: you need multiple indeces. Cooper also points out the user interface principle of enabling the user to perform the same tasks in different ways, a concept similar to having multiple indeces as they allow you multiple approach angles to the same piece(s) of information.

Bush takes the concept of retrieval one step further, instead of using the familiar indeces his idea is to allow information to be stored by association. This is an idea which falls naturally for humans as we often store something by associating it with something else we already know. To find something then means you do not necessarily need to know specifics, instead you can know that it's related to something else, say that it's preceeded by another article for which you already know the author and title of. As you find the specific article the association to the one you're looking for allows you to find it, in a similar fashion to remembering things by associating them with parts of a story. Bush's idea also allows us to walk through information in a similar way to walking on paths through a forest, which is a key concept of the system he describes, it is called the Memex.

2.1.3 The Memex

Unlike many other ideas, Bush not only described the concepts behind the Memex but also saw it as a physical object; the scientist's desk. His vision was an upgraded version of what we all know, adding devices for input and output. Images of documents, objects and anything else was to captured onto microfilm, the most advanced storage technology at the time, so the desk had a camera for that. There was also a keypad for entering queries, and two monitors for looking at the resources that were available. In addition he describes capturing audio and video to enable the scientist to work freely with her hands. This makes the Memex not just about the concepts and the physical object, it's also an idea about the scientist's workday as a whole, a holistic approach we will see again when we come to Douglas Engelbart in section 2.2.

At the core of the Memex is the idea that information can be stored by associating two items to each other, tieing them together. This means that as we are working, looking up information, reading, creating notes, we can create what we deem a natural connection between two pieces of knowledge. As the day progresses we keep connecting pieces to each other, creating what Bush describes as a “trail”, a path through the forest of information. This path can then be stored so that we later can consult, extend and store it again, in line with the earlier principle of usefulness.

Bush also understood that scientists do not exclusively work alone, and “As We May Think” also describes how these paths can be used in a collaborative fashion. Two scientists talking about a subject can discover that one of them has information available in form of an already created path, and this path can be transferred to the colleague’s system through an export/import process. This exchange of paths means someone else can pick up where you left off and extend it, but they can also walk through and review what you have done. A key feature could be the possibility to learn knowledge through reviewing other’s work, making sure it makes sense, and perhaps adding notes along the way, or splitting off to explore a different path. Vannevar Bush’s idea about the Memex is thereby not just an abstract concept, not just a physical object, but a way of *working with* information. In today’s world with the proliferation of web applications, it is easy to see that we can extend these ideas to create a collaborative web service, enabling people to work together on a path or a set of paths, and we’ll return to the trails in more detail in section 4.5.

2.2 Augmenting Human Intellect

Douglas Engelbart is perhaps most known for inventing the computer mouse, which he did in the mid-sixties whilke working as a researcher at the Augmention Research Center (ARC). The ARC was a research lab that he founded while working at the Stanford Research Institute, and as the lab’s name signifies, their projects revolved around concepts that we will delve into. Another thing that Engelbart is famous for is what has become known as “The Mother of All Demos”, his demonstration of the On-Line System (NLS) which took place in 1968 at the San Francisco Civic Auditorium.

In the sixties and mid-seventies, Engelbart was on the forefront of computing research, augmentation and NLS being two main subjects that created a lot of interest. One might say he was before his time as he has since whiffered off into obscurity, his research papers are largely unknown, shadowed by his invention of the mouse. In our case it is high time to bring him back as he builds on what Vannevar Bush wrote and provides insight and descriptions of ideas that take Bush’ concepts to the next level. It’s 1962, Engelbart’s article is called "Augmenting Human Intellect"[8], and as we'll see it is an article with important ideas.

2.2.1 Augmenting what exactly?

Augmenting your intellect probably sounds tempting, but what exactly does it mean? At first thought you might think it means the ability to increase your intelligence, boost your IQ, make us all able to join Mensa[1]. Engelbart points out that this is not an incorrect interpretation of it, but that he would describe it differently. He seems to have a more practical approach to what it is, and points to the effects of augmentation in his description:

> "By 'augmenting human intellect' we mean increasing the capability of a man to approach a complex problem situation, to gain comprehension to suit his particular needs, and to derive solutions to problems."

This brings us back to Vannevar Bush because he was also concerned with much the same kind of problems, except that Bush wrote his article at a time where computers were in their infancy, and he therefore did not envision a computerized solution. Regardless of that they do talk about the same thing, looking for a technological solution that can enhance our ability to work with and solve challenging problems. There's also an ambiguity in that, as it can mean both solving problems at a higher rate as well as being able to solve more complex problems. We're now quite skilled at getting higher performance out of computers, both by inventing faster algorithms and by creating bigger and faster computers, but as far as we understand Engelbart, he was not primarily concerned with those problems that can be solved simply by throwing a bigger hammer at it. Complexity in his realm meant solving those complex issues that are not readily suited for automated solutions.

Where Bush asked what kind of challenges would be worthy of our best work, Engelbart writes that he is seeing augmentation as a means of helping us understand the complex issues in society. He does not specifically name any, but we believe he is referring to some of the complex issues that surround us ever day, like poverty, obesity, environmental issues, and epidemics like AIDS and malaria. Could we create a system that makes it easier to understand and tackle these seemingly unsolvable problems, perhaps figure them out once and for all? Those are the kind of questions Engelbart seems to be interested in.

[1] The barrier of entry is scoring in the 98^{th} percentile. Whether IQ boosts would actually get you in is left as an exercise for the reader.

It is also evident in his writings that he does not see technology as an isolated solution. Instead he points out that being successful in getting this to work means we not only change what we do in that we use the technology, it will also affect a lot, if not all of everything else too. As we saw earlier, the Memex affects not only the scientist in her everyday work, it also affects colleagues as they will have to adapt the same system to be able to exchange data. Engelbart also sees this, he describes the processes that go on around us all the time and that they'll also need changing, again a holistic approach to the effect of the technology. You might say that he sees the introduction of the kind of technology and processes described in "Augmenting Human Intellect" as the kind of complex problem that he is concerned with solving, and that his ideas also go towards helping to create suitable technology and introduce it in a way that will make it work.

2.2.2 Symbolic Structures

Engelbart's key concept was to utilize symbolic structures as a way of making us able to store the thought structure of our mind. One examples he uses is his system of notes, which he writes during his research and then indexes using a complex system. He sees that this solution is becoming archaic and is looking for a better way. The symbolic structures he describes are similar to the associative structure that Vannevar Bush wrote about, the main difference being that Engelbart describes them in a less specific manner. A more traditional approach to the issue would be using a linear structure of words to describe the structure, meaning a story, or a thesis like this one. The question is then if it would be easier to get your point across if you could map everything out? An illustration of what a part of the map for this thesis could look like is shown in fig. 1 on the following page.

Although somewhat less specific than Bush, Engelbart saw the symbolic structures being similar to the associative trails described in "As We May Think". Engelbart describes the symbolic structures as a way of breaking down a problem so that a human can walk through it. You can look at it as akin to how we describe a problem to someone else, we need to break it down and structure it. The main difference between a static structure and communicating with someone else is of course that they have the ability to ask questions, changing our description, so the structure we start out with might not be the one we end up with.

Just as collaboration was not absent in "As We May Think" it is easy to find in Engelbart's work. The On-Line System he showed in "The Mother of All Demos" as we mentioned earlier was already a networked collaborative system. Engelbart also had the holistic approach to it all, seeing that the computer system would not be separate from society but working as a part of it. Aiding one person's way of working would mean it'll affect everyone's way

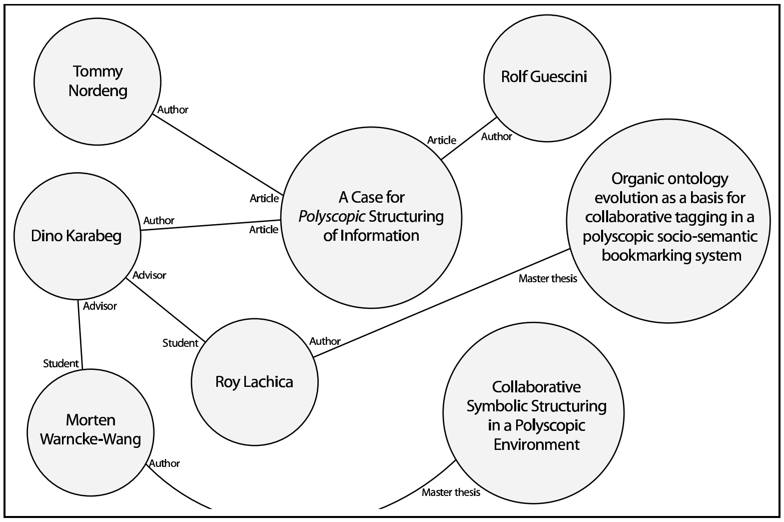

Figure 1: Map of a part of this thesis

of working, and to be able to solve the complex issues we mentioned would mean enabling people to work together on the structures. Engelbart clearly understood that work all over the globe is not about individuals working alone, it needs to support both the groups working together as well as being able to map an individual's mind.

2.2.3 No revolution

One really interesting argument made in "Augmenting Human Intellect" is that for the ideas to be succesful they cannot happen as a revolution. In Engelbart's view people rarely work that way, instead they work in small steps, just like breaking down a problem into smaller pieces as we referred to earlier. A way to look at it is to imagine moving a million people a thousand kilometres[2]. Doing it by revolution would mean you move them all in an instant, while the other approach is to get them there one step at a time. The latter will take quite some time, the former might mean you leave people behind, or it could never work at all. In "The Logic of Life"[13], Tim Harford takes a rational approach to revolutions and point out how it is a question of swapping the power in the revolution for a lasting future change. History tells us there are several examples of revolutions not working out in the long run, and it could be that Engelbart had that in mind when he asked us to get where we want one small step at a time.

[2] Which is approximately 612.37 miles for those who prefer imperial units.

2.3 Creating *Polyscopic* Information

While none of the first two pieces of my foundation are less than 40 years old, the third piece is just a couple of years old. Written in 2006, "A Case for *Polyscopic* Structuring of Information"[12] takes a modern approach to how we go about creating information, marrying the methodology of Polyscopy created by Dino Karabeg with the technology of Topic Maps. Although contemporary and radical in some ways, as we will see it is not orthogonal to the other two. Guescini et al also refer to the problem of handling the vast amount of information available to us, and bring a new idea to the table, aiming to turn the information we make into something that is better suited for us.

2.3.1 The Principles of *Polyscopy*

At the core of Polyscopy is the *scope*, and the name of the methodology literally means "multiple scopes". The scope is defined as "the point of view" and can be compared to looking through a telescope or a pair of binoculars to examine something far away. In those circumstances the view is limited, and in Polyscopy the scope is limited too. It is dependant on the environment in which we work, meaning the choice of subject, the language we use, and the culture in which we write, where for instance the scientific tradition is perhaps most relevant. These choices determine what we are able to express, thereby also determining what we are able to see, in much the same way that our vision is limited in that we cannot both look at something afar and something close at the same time. The methodology of Polyscopy asks us to keep this limitation in the information we produce, where a scope is defined as *coherent* if it corresponds to a single point of view. By that they mean that we do not attempt to see both the whole forest, a single tree, and a small ant on that tree at the same time. Instead we are to describe each phenomenon in their own scope, make sure each scope is coherent, and then provide the reader with the ability to choose the scope that corresponds to what they are looking for. Thus we create multiple scopes, from which the name "Polyscopy" derives its meaning.

The second principle is that of abstraction. To be able to better utilize information we must split it up into manageable pieces. We already do this in writing as we create sentences, paragraphs and chapters, and use lists and tables. Polyscopy asks that we not only continue with these abstractions but also make sure that they are connected, because central to the principle of abstraction is the idea that it enables a way of navigating through information, much like you navigate your way towards a destination with a map.

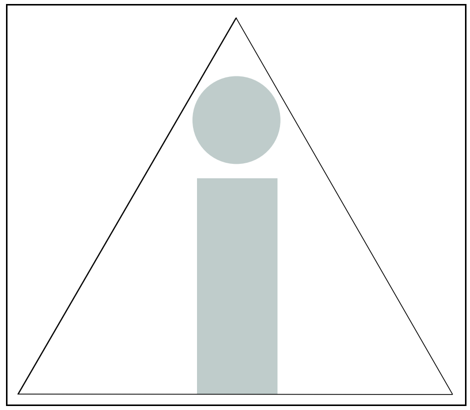

Figure 2: The “i” ideogram

Polyscopy defines three types of abstraction: vertical, horizontal, and structural. *Vertical* abstraction means creating a structure that enables changes in level of detail. A metaphor to use for describing this is looking at our body of information as a mountain, rising up from the savannah like Kilimanjaro in Africa. From the top of the mountain we can get the overview, see the bigger picture. We can look far around us, but are unable to examine the details in the ground below. The opposite is found by moving to the bottom of the mountain where we can inspect the intricate details but we cannot see clearly what is found at the top. These kinds of views are what Polyscopy's vertical abstraction principle as us to create, as many levels as need be to allow movement from overview to fine detail.

Another way to look at it is the “i” ideogram shown in fig. 2, also symbolizing the vertical abstraction. The top is the dot of the i, where we find the overview. The round shape signifies the removal of detail of provide us with an understanding of the bigger picture. Underneath the dot is the stem, the structure that supports the dot, having increasing level of detail as we move downwards. Surrounding the i is the triangle, resembling the mountain of information as we described in the previous paragraph.

Secondly there's the *horizontal* abstraction, which defines that instead of lumping all information on a specific level of detail into one big mess we are to separate it into smaller, comprenhensible pieces of information. Again this is similar to what we already do in that we create paragraphs and chapters, but we must remember that each abstracted scope should correspond to a single point of view as stated earlier. The final abstraction is *structural*, where the point is that the structure created by splitting up horizontally vertically is one that we must be able to reconnect. The separate pieces must have a relationship such that we can recreate the whole picture, rather than stay independent.

This brings us to the last principle, that of the *perspective*. Where the other principles define how to organize information, the perspective principle is all about what information we seek and provide:

> "The *perspective* as criterion directs us to seek the information that we are lacking in order to understand the whole, and to present that information in a way which makes the character of the whole and the relevance and the roles of the details clear."

We are to go look for all that is available, bring our spotlight with us and illuminate that which is in the dark, rather than show that which is already visible. Again a holistic approach can be seen, as we are expected to cover all sides of an issue, also those that might oppose our primary standpoint. The reader should get a clear and concise picture of the whole situation, and thus make them able to come to their own conclusions.

2.3.2 Motivation

Just like Bush and Engelbart before them, Guescini et al anchors part of their motivation in the problem of information overload. They point out that the amount of information available and the way it is organized does not help us solve the really difficult challenges our society faces. Topics like declining health and non-sustainability are mentioned, where having to work your way through an ever growing jungle of information is not in any kind of way a help towards a solution. Another way to look at it would be to envision information as islands in the ocean, and there are no bridges between any of them. The structural abstraction we saw in the previous section are then the solution to this isolation, recreating the bridges between the islands. This will allow us to move freely between them in a similar way to how the paths of Bush and Engelbart allows us to find our way as trail directions in a forest.

The final piece of motivation is the idea of *designing* information, as opposed to unconsciously creating something we believe might work. A keyword here is "tradition", that what learn and follow the traditions of those before and around us, without necessarily questining them as much as we might need to. Again we reflect back to the topic of information overload, and it is quite evident that we're not yet at a place where we have the technology that makes us able to sort through it, otherwise there would be no such thing as "information overload". As we saw in section 2.1, Vannevar Bush raised concerns about information magnitude and the need for specialization already in 1945, and Guescini et al argue that instead of attempting to once again find a technological solution to the issue, we are to make conscious choices about how we go about creating the information. We should *design* the information. Build it using a structure that enables easier grasp of the topic at hand, as well as a way of understanding the relationship between the piece of the structure to make it easy to not only understand the general concept but also the intricate details. Do it in a way that covers all aspects, also those that which are mainly left lingering in the dark. Leave tradition behind and design information for the future.

2.4 The World Wide Web

As we come to the fourth pillar in our foundation we move to a specific piece of technology, the World Wide Web. It is perhaps technology that is taken as a given nowadays, in a world where we the cutting edge has arrived at Rich Internet Applications (RAIs) like Google Docs[21] where we can collaborate on documents and spreadsheets, and presentation tools like 280 Slides[16], both providing services usually found in standalone applications but delivered over the Web. There is also "widgets", small applications often run either in your browser or on your desktop to provide useful tools like stock quotes, weather reports, or a summary of the latest news, maybe combined in ways you couldn't have imagined. Lastly there's "mashups", the principle of combining available Web-based services in a way that's useful, like HousingMaps[33] where you can see apartments or houses from Craigslist ads display on a map from Google Maps. A similar example is geoGreeting[44] which utilizes images of buildings shaped like letters to write a greeting you can send to people. These kinds of services were perhaps not what Sir Tim Berners-Lee had in mind when he created the Web, as we can read in "Weaving the Web"[1] it started out with something a lot less complicated.

The World Wide Web started out as a program called "Enquire", which Berners-Lee named after a book first published in 1856 called "Enquire Within Upon Everything". This book attempts to answer questions about nearly everything, and the program Berners-Lee created in 1980 while working at CERN[3] similarly attempted to answer all questions he had about the organization he was a part of. He describes the work situation at CERN as chaotic, particularly for newcomers, as the organization was fluid with projects group created ad hoc after engineers talked over coffee in the hallways, and disbanded when funds were unavailable. This made it difficult to keep track of who worked in the various project groups, what software they created, what computer resources they needed and used, what documentation was written, and most importantly, who to call in an emergency. Berners-Lee's Enquire program utilized bidirectional – two way – links between internal objects to represent the connections in the organizational puzzle, enabling him to map everything and query it to get answers, a structure similar to those described by Bush and Engelbart. The program could

[3] The European Organization for Nuclear Research. Its abbreviation comes form Conseil Européen pour la Recherche Nucléaire, the original French name.

also link to external objects, but would then only have a unidirectional – one way – reference. He writes that he chose this approach because bidirectional links would require storing the same information in two places, which sooner or later would mean they get out of sync, and that storing and displaying the incoming links could mean you would have to manage and show maybe hundreds or thousands of links. When Berners-Lee left CERN because his consulting time was up he also left Enquire behind, and the disk with it was later lost.

Berners-Lee returned to CERN in 1984 and once again thought about the Enquire program and what could come out of it. He envisioned a larger system able to map any kind of information, not just the information at CERN. This time around he did not create the bidirectional links, instead he started with the unidirectional as it was easier to program. We have already seen a couple of drawbacks with the bidirectional links, but their benefit is that you get to know exactly who links to you, while on today's Web some search engines can show you who might currently be linking to you. Getting to know who refers to you was not as important to Berners-Lee as the decentralized structure that would come from having unidirectional links. It enables you to refer to someone else without having to ask their permission first, there is no central authority you need to clear it with. When you combine this with a system where you are free to create and edit resources you facilitate rapid growth which is what Berners-Lee had in mind. He though add the bidirectional links in back later, but as we know by now, they're still missing from the Web. With the addition of a NeXT computer on his desk he found it easy to create a working prototype of the system based on these ideas, consisting of an application to edit and access documents, the "browser", and a server program that enables queries for documents to be delivered. As we see there were very few rules and principles behind the first prototype of the Web, and from what Berners-Lee writes this was a conscious decision:

> "I would have to create a system with common rules that would be acceptable to everyone. This meant as close as possible to no rules at all."

This principle was kept alive throughout the development of the Web, even as the year passed with it being nothing but an interesting and promising idea. In the late nineties, after the Web was successful, Berners-Lee moved to Boston to help run the World Wide Web Consortium (W3C)[45], an organization set up to develop and maintain web standards. It takes care of central Web technologies like the HyperText Transfer Protocol (HTTP), HyperText Markup Language (HTML), eXtensible Markup Language (XML) and Cascading Style Sheets (CSS). Keeping with the decentralized principle mentioned earlier, the standards are not issued as standards in the traditional sence, instead they are *recommendations*, describing what the members of the working group behind each recommendation believes is the best way to do it. You are of course free to publish information adhering to the standards, or you can publish them in a way you prefer yourself. There's no central organization telling you what to do, and many new standards start out as an idea that people pick up on, like MicroFormats, which are non-standard additions to (X)HTML used to solve minor needs.

One of the ideas of Berners-Lee that could be argued has been fairly absent from the Web is that it was to be about writing, just as much as it is about reading:

> "My vision was a system in which sharing what you knew or thought should be as easy as learning what someone else knew."

As we mentioned earlier, the prototype that Berners-Lee created had a browser that was also an editor, making it easy to create and edit documents, not just access what others had shared. Very few browsers today have this duality, the most well-known is probably the W3C's reference implementation browser Amaya. Publishing was at first cumbersome for those not into learning the technology, which lead to the creation of websites like GeoCities[4] where everyone could publish with the help of templates and easy-to-use tools. In the mid-to-late nineties Usenet, a discussion forum where "newsgroups" are each dedicated to a different subject, was also a competitor to the Web, but because Usenet is based on the Network News Transfer Protocol (NNTP) we do not regard it as Web technology, and to access it meant using a particular kind of software called a news reader[5].

The late nineties saw the start of "weblogs", now simply called "blogs". Nowadays it is a template-based content management system geared towards publishing regularly, daily or perhaps multiple times a day, and this makes it easy to start publishing. With a public blog service like Blogspot or Blogger.com, you could have your first post out in about five minutes. This kind of easiness makes us get a lot closer to Berners-Lee's goal, although the browser is still regarded as a browser and not so much an editor (with today's Rich Internet Applications and scripting that line is definitively blurred). We've slowly made our way towards having a Web where sharing and reading are both easy, but does that mean we contribute? Research into the matter brings back results indicating that a disproportionate amount of people are the ones contributing. Jakob Nielsen indicates a 1/9/90 distribution[29], 1%

[4] Which was bought by Yahoo! Inc in January 1999 and is now called Yahoo! GeoCities.[26]

[5] Web -based gateways have been created, the most famous one is maybe DejaNews[48], bought by Google in 2001 and is now known as Google Groups[28].

contributes most, 9% contribute irregularly, and the last 90% never contribute at all. In a study of contributions to Wikipedia, Priedhorsky et al indicate that 0.1% of the editors contribute the 44% most seen content, with the top 10% of editors contributing 86%[35].

The big question is of course then whether the Web is about information and sharing. Priedhorsky et al's results are interesting, but they refer to the content which is most seen, most popular. Another aspect of the Web is what is now called The Long Tail, which is mentioned by Tim O'Reilly who we will get to in section 2.5. The Long Tail is all about the smaller sites, each one not getting as many visitors, but when you combine them it amounts to huge numbers. One result of the Long Tail is that small bookstores around the US can keep open partly because they keep selling books through Amazon.com's reseller program. David Weinberger, author of "Everything is miscellaneous" argues that the Web is not about information, it is about communication[22], which brings us back to Tim Berners-Lee because he also has an alternative view on the Web:

> "The Web is more a social creation than a technical one. I designed for a social effect – to help people work together – and not as a technical toy."

This reconnects us to the main idea behind this thesis, collaborative creation of information resources, groups of people working together to share knowledge, hoping that this enables deeper understanding of complex problems, going back to the ideas of both Vannevar Bush and Douglas Engelbart.

It is not a primary concern of this thesis, but still worth mentioning, the visions for the Web that Tim Berners-Lee puts forth at the end of his book. He dubs it "The Semantic Web" and it is a Web where technologies for adding semantic information, which is information describing the actual meaning of the content you publish. Berners-Lee sees this Web as having the ability to make computer do more complex computation based on meaning, rather than just add things up. One example similar to those described in his book is finding the answer to questions like "When can I get an appointment with my G.P.?" In the semantic Web, the computer understand the intricasies of this question, search your G.P.'s calendar to find an appointment that fits with your calendar, and could perhaps also provide you with driving directions and pay for parking and any road tolls on the way in advance. The first thought might be that XML is the solution to this problem, but although XML allows easy exchange of data between computers it does not come with a ready set of specifications for describing meaning. Instead the W3C has a specification using XML to create a language for semantics called Resource Description Framework (RDF), which they hope will add a layer to the Web that enables these computerized solutions. As Berners-Lee puts it, a Web where human thinking can be combined with the power of the computer:

> "...the Web will be a place where the whim of a human being and the reasoning of a machine coexist in an ideal, powerful mixture."

This thought brings us back to Bush and Engelbart's idea that we should seek ways to automate tasks so that computers can do them. Rather than have us become bored with monotonous tasks the computers can do it and leave us to spend our time working on the interesting and difficult problems.

2.5 The Second Generation of the World Wide Web

In 2003 the term "Web 2.0" began popping up as a descriptive term for a new breed of popular websites. How exactly were these different from the previous generation of websites? The answer to this question is the fifth and final pillar of our foundation, found in Tim O'Reilly's article "What Is Web 2.0"[32], where he explains the key differences in the design patterns and business models of the new companies.

2.5.1 The web as a platform

The Web as initially created by Tim Berners-Lee is often seen as being about publishing documents for others to read. For many years, Web pages were mostly static or driven by complex technologies and developing these kind of services was a complicated effort. As Tim O'Reilly points out, at that time the business side of the Web was about selling systems and servers. The company Netscape sold both their web browser, Navigator, as well as a web server, and aimed to keep users locked to their systems, their platform.

This changed as the century drew to a close. Netscape's web server was surpassed by Apache, which has now been the most used server software for the past 12 years according to Netcraft[25]. In combination with freely available open source tools like MySQL and Perl/PHP running on top of a Linux operating system, those looking to get into the web business had a low cost solution available. You could of course choose other platforms and programming languages too, Sun's Java was available through Java servlets and JSP, and on the Microsoft platform you had ASP. The basic needs were covered, developers were not at all locked in to any platform specific solution, and so new ways to attract users were needed.

Tim O'Reilly describes this as value moving "up the stack". Instead of delivering platforms we moved to delivering services over the web platform, as the platform itself had matured and become transparent. This also meant that it was a lot more difficult to lock an application to a specific platform. The way Netscape attempted to do this was to deliver both the web server and the browser thereby controlling the user, but as they lost the browser war to Microsoft the user was outside of their control. Today the picture is even more complex as users view content using any kind of device be it a desktop computer, a set-top computer for their TV, their mobile phone, or even a gaming device as both the Nintendo Wii and DS are web capable. Companies now need to play to the strengths of the Web platform, and we'll examine three of those strengths as mentioned by O'Reilly.

2.5.2 The user is the key

What could be the most important difference between the first and second generations of websites is how they treat their users. In the first generation, users were customers similar to those walking into a store, and from the website owner's point of view they were there to generate cashflow, either through buying items or services, or by seeing advertisements. There is a difference in the ads too, they have gone from just the traditional billboard-/print ads delivered by companies like DoubleClick to also being the tailored, text-only ads delivered by Google and Yahoo!. The ads are also utilizing the power of the Long Tail, nearly everyone can have ads on their websites and each earn a few bucks and when combining them all they amount to a large number of ads.

Secondly there has been a move from seeing visitors as customers to seeing them also as producers. Amazon.com is a prime example as they quickly allowed users to write reviews, and now also have them add keywords to the description of products, and make lists of products they think other people should read. Another important change done by Amazon.com was to have user behaviour in itself become important. They let things like what you search for, what you buy, and what's on your wish list be used for suggestions for what you might be interested in buying next. The latter is what we call a "recommender system"[3], and we will come back to them in section 3.2.1. Tim O'Reilly calls this principle the "architecture of participation". Your users are a resource, and if you treat them as such they can add value to your system. As we'll see next, one way is to make your data worth more than your competitors.

2.5.3 Your data is an asset

A data resource is easily replicated if there are multiple distributors of the same resource. The example O'Reilly uses is that of map data where there are several companies you can get that kind of data from, and they mostly deliver the same kind of data. Through user participation those data can be extended, and Google Maps is a good example. Google supplies an API that allows the users to manipulate the maps, and the result of this are the kind of mashups mentioned in section 2.4, like HousingMaps and geoGreeting. Amazon.com is also one example of this kind. Looking at their database of books, it is not unique in and of itself, but as users add reviews, keywords and create lists of books they add value to it. This makes the barrier to entry into the market higher, so it's difficult for others to outcompete you.

2.5.4 Think syndication, not coordination

The traditional approach to owning an information resource is to restrict access to it, perhaps as a protection from it being used in ways you did not intend. Web 2.0 goes in a slightly different direction, their principle is to create services where information is exchanged and syndicated. This allows users to build new services upon your data, or intergrate your information into their site adding value to their site, and they display that you allow this kind openness. The mashups described earlier are again an example, for instance HousingMaps combines two available services, Google Maps and Craigslist, to create a third service which also is useful. Attempting to create, coordinate and maintain all these kinds of services through your own company is probably impossible, which is why syndication is the way to go. O'Reilly's design patterns are also found when we look at the business side of computing, as trends described by Standards & Poor's[2] are that companies go for delivering services, and utilizing the business model of open source software where it is an advantage, thus opening up rather than keeping their information behind lock and key.

3 The Sketch – Connecting the bigger dots

Having built a foundation from our five articles it is time to look at some of the ideas they put forward, and start moving towards making these a reality. We are not yet looking at the implementation, instead we will start sketching some general ideas and see what kind of issues arise from that and how they can be solved. We will also bring in some of the ideas necessary to move from the conceptual descriptions found in our articles towards actual implementations. Let us start with a look at social software, as we're trying to create a collaborative tool.

3.1 Social software/CSCW

The field of Computer Supported Cooperative Work, or CSCW, is not new, people have been using computers to work together for several decades, one example was Engelbart's On-Line System. As we discussed in section 2.4 when we talked about the World Wide Web, it was for long regarded as a place for publishing information rather than a place for collaborative work. This has changed drastically in the last few years as what is known as "social software" has become a household term. If we look at the list of the most popular websites in the world as according to Alexa[17], three of the top 10 sites are clearly social sites; Myspace (6th), Wikipedia (7th) and Facebook (8th). Although not primarily regarded as a social site, Yahoo! (1st) and YouTube (3rd) have also a strong social element present, they're not respectively a simple search engine and a video viewing site. Number 9 on the list is blogger.com, which is a freely available blogging tool, and there is a lot of social interaction there too for instance as users comment on each others posts and link to each other.

Tim O'Reilly was also talking about social software in his article about web 2.0 that we discussed in section 2.5, in that it is an aspect of leveraging your users and adding value to your website through their creation of content, perhaps even by them simply interacting with each other. When it comes to examples of social software the most popular one is perhaps Wikipedia. It's the 7th most popular website at the moment which is quite the feat considering that it's a collaboratively created encyclopedia. The exciting thing is the main idea behind Wikipedia, that anyone can edit an article, but you can choose to register if you like. Research has shown that it contains a similar amount of errors as a reputable encyclopedia like the Encyclopædia Britannica[52, 24]. One of the advantages of an online encyclopedia like Wikipedia was that they could simply make a list of all errors and start correcting them, and little over a month later they were all corrected[53]. They also made sure to correct the errors found in the Encyclopædia Britannica[51] too.

One of the major ideas behind Wikipedia is of course that they aim to be an encyclopedia, and this leads to the ideal that their articles should be unbiased. In Wikipedia terms this is referred to as NPOV – Neutral Point of View. It is an ideal that can be difficult to reach, and each entry in Wikipedia has a discussion page where readers can give feedback and editors can discuss issues. When editors disagree they might end up is what is called an "edit war"[50], a heated debate where they move an entry back and forth between different revisions. In "History Flow"[43] they display this problem graphically, and the question is how can the issue be resolved? To begin with Wikipedia had few mechanisms and policies for these problems, and over the years they have developed them and documented the best practices in their list of policies[54]. An additional mechanism that is powerful in Wikipedia is the ability to monitor pages. As a registered user you are able to receive email updates when changes are made to an entry, and that this works for popular pages is evident by the quick corrections of acts of vandalism. Viégas et al report median times for deletion of the contents on a page to 2.8 minutes, while the median for all content being 90.4 minutes. Either way it is resolved fairly quickly.

At the other end of the social software spectrum we have the sites who are mainly about social interaction like Myspace and Facebook. Both have secondary services, Myspace having video upload, a popular music site and blogs/diaries, while Facebook has applications, user created software services for Facebook that interact third party services if necessary. Central to many of these sites is that they allow users to map out their social network by befriending other users. Compared to our dynamic relationships when communicating with each other, shown for instance by how a simple question like "Where are you?" asked by the same person will be answered differently depending on the situation[15], the idea of what a friend is on these sites can be described as very static, on both Myspace and Facebook friendship is mainly a binary unit, you're either my friend or you're not.

We will return to the social side of the Web in section 4.4, and hopefully show how a combination of some key features can keep the strengths while making some of the drawbacks less of a problem.

3.2 Information overload

By *information overload* we mean the condition where you are unable to make decisions or stay informed simply because there is too much information available[27, 11]. You feel overwhelmed and spend all your time running around between different information resources as you try to get to know more and more as you're unable to make a decision until you know all there is to know about something. Given the rapid rate of expansion in the information space the conclusion is you'll never catch up, for instance on YouTube approx. ten hours of video is uploaded every minute[42].

Information overload is a real problem, and as we've seen it's a concern of both Bush and Engelbart, and it was also a motivation for the creation of the Polyscopic methodology. Additionally we see that it is easy to get lost on the Internet because the possibilities are so many, as Barbara Ehrenreich writes in "Baith and Switch":

> "The time I spend on the web has a dank and claustrophobic feel. After traversing a few links, I forget where I started and am lost among the pages full of advice, support groups, networking events, and coaching opportunities geared to various salary levels."[7]

This is not a psychology thesis, so we have chosen to not attempt to solve the problem from that perspective. Instead we will be looking at approaches in information technology to make the amount of information less of a burden. First we will tackle it from the output side, seeing if we can do something about the amount of noise in the existing pool of information by filtering. Secondly we will look at the input side to see if we can alter the quality of the information we're creating.

3.2.1 Social Navigation

Social navigation in the real world is akin to following people around when you're looking for the exit at a subway station, or choosing to eat at a restaurant or go to a concert because lots of other people are there. It might not always be successful but it's a good bet. In computing social navigation is not so much an old discipline, one of the first examples referenced comes from the early nineties when the Tapestry system was created at Xerox PARC[3, 10] where user relationships helped determine if an email was worth reading. You told the system whom you trusted and it would then highlight emails that they had deemed readworthy. This is the start of what would later be known as *recommender systems*. As we've seen earlier, Amazon.com utilizes this kind of system when it automatically recommends products to you that match those bought by others. There are also several other approaches to filtering out useless content and we'll start with looking at Usenet again.

Usenet is a network of servers allowing users to discuss topics in what is called "newsgropups" as we first mentioned back in section 2.4. When the Internet became popular in the mid-nineties the signal to noise ratio of many groups fell as the amount of participants became too large, and partly also because some users found that Usenet was a cheap way to advertise. With the number of daily postings in a group reaching several hundreds it's difficult to keep up and that is where the recommender system comes in. GroupLens[37] was a system enabling participants to rate Usenet postings based on their quality, and you could then choose to skip articles with low ratings because those who had seen them before you declared them as useless. Thus a social filter is created and you're able to utilize the wisdom of the crowd.

Today there are two popular websites that work in a similar fashion, Digg[19] and Reddit[18]. Users submit stories they believe are interesting to a queue, and the other users can then give a thumbs up or thumbs down, in Digg lingo they're called "digg" and "bury" respectively. If a story gets enough votes it ends up on the front page of the site. Similarly to how GroupLens worked these sites enable you to let the collective intelligence of those before you decide what is worth reading.

These sites do not come without their drawbacks. Digg was without categories until version 3 was launched a year and a half after the site went public[49], and it is easy to see the lack of categories as a problem. If you're not interested in politics while the majority of Digg's readers are, and you cannot remove the whole category of politics, you'll have to wade through a bunch of stories that are completely uninteresting. When your system is supposed to help filter out those kind of stories, it's not working right if you have to keep skipping through them. This was never a problem in GroupLens because Usenet newsgroups by nature are restricted, people do not mainly discuss politics in a Linux-related newsgroup like comp.os.linux.misc.

The second problem area is whether it is possible to gang up and game the system. There was criticism against Digg because people believed a group of users rigged the site by burying a story, which means they vote them down, rather than just abstain from voting. Digg answered the challenge by working on its algorithms to help tackle the problem, although they do not disclose how exactly their algorithms work because they don't want users to be able to figure out ways of gaming it[20]. Basically we're looking at a problem where the individual's preferences do not match those of one or all groups, and the question is then if it is possible to design a system where the needs of the group and the needs of the individual are balanced out. If the user is mostly interested in a niche we should enable her to pursue that interest, but also benefit from the intelligence of the group where possible.

3.2.2 Polyscopic solutions

As we pointed out in section 2.3.2, one of the motivations behind Polyscopy is information overload, and the suggested solution is to make the information we create better by *designing* it. Social navigation only helps to remove noise from the existing corpus, it works on that which is already created.

Polyscopy takes the alternate approach in that it states that we should make conscious design choices and go against tradition if need be, to make sure that the information we create is already in a form that is more useful to us. It means we are to imporve that which we put into the corpus, rather than spend resources on how to better take stuff out of it.

The vision of this thesis is to create a system that combines these approaches, enabling users to make information more useful by changing the way it is created as well as changing the way we organize it and find what we're looking for. This should not just work for the individual but also for smaller and larger groups as we're certain that this kind of interaction with the crowd is a way to augment our intellect.

3.3 Information Design in three dimensions

3.3.1 Three dimensions of Polyscopy

Section 2.3 gave us the three principles of the Polyscopic modeling methodology; the *scope*, the *abstraction* and the *perspective*. As you might remember the scope is our point of view, and because we are to limit ourself to *coherent* scopes we use abstractions to split things up into comprehensible pieces. There is three types of abstractions used in Polyscopy; vertical, horizontal and structural. Lastly we have the *perspective* which stated that we are to uncover the information that is hidden in the dark to reveal the *whole*.

Another idea found in Polyscopy is that the structure created by abstraction is to be three-dimensional. At first it might seem that we're simply working in two dimensions as there are variations in level of detail as we move up/-down the vertical axis, and between different scopes as we move along the horizontal axis. In a class project for the course INF3210 – Information Design here at the University of Oslo, a JavaScript-based polyscopic navigator was created[46]. It lets the user navigate a two-dimensional polyscopic topic map to learn more about the XSL-FO technology. Because it was based on a small information structure that was already laid out as a tree, staying with that 2D approach was appropriate, but it leaves us with the question of what kind of conditions need to be met for the three-dimensional information mountain to appear. To illustrate the issue we will use a couple of examples, starting with the simple example found in fig. 3 on the following page.

We're far away from the *large and complex* issues described by Bush and Engelbart, but even at this small scale we can identify a characteristic problem, how to decide what level in the structure something is on. Looking at the structure in fig. 3 on the next page we can argue that the main subject, aptly named "Main Subject", is the top level. From that assumption it would

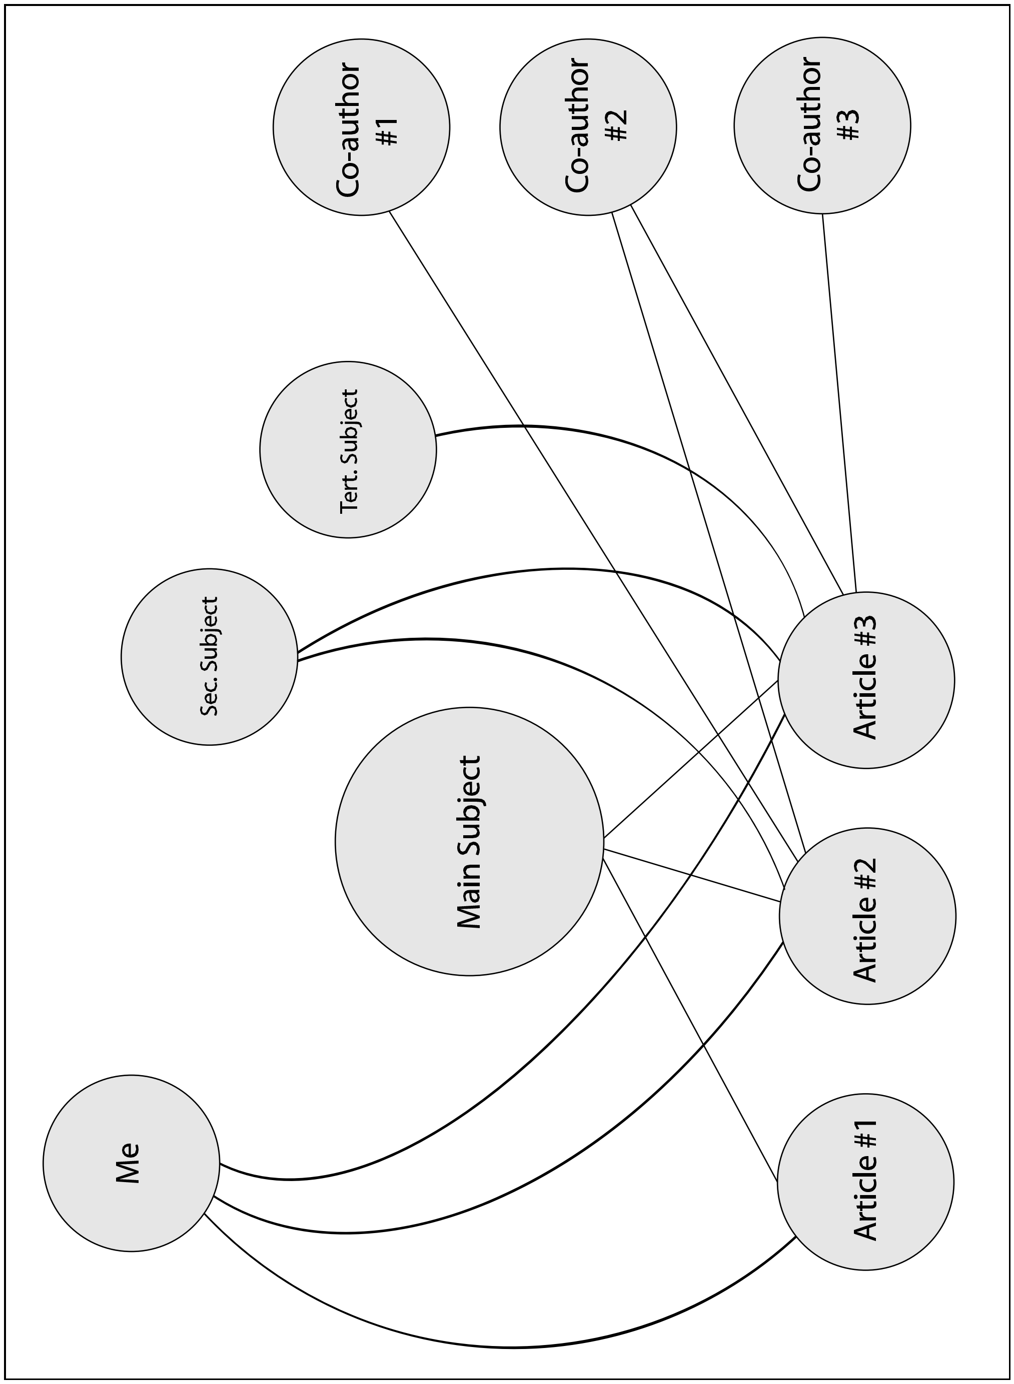

Figure 3: Example symbolic structure from scientific articles

follow that the three articles are on the second level and that the authors and the minor subjects are on a third. If we instead approach this structure from my point of view, "Me" will be the top level, and the other levels will follow from there. If you're one of my co-authors you might disagree and prefer that you be the top level. The result is that we have many approach angles to the same structure and what is right for you might be wrong for me, and vice versa.

Instead of having to spend time deciding exactly where everything fits and what level it is on, our approach is to look at the structure as one whole piece. If we have read the three articles and written a summary for others to learn from and want to connect this summary to the structure, we connect it to the *whole*. This means we see the structure in fig. 3 on the facing page as a complete structure describing what we currently see as the context for our summary. Thus our key insight from reading these articles becomes a level above the structure, and we have gone from a two-dimensional approach to a three-dimensional one.

3.3.2 Structural User Interfaces

Keeping with the principles of Poylscopy means we should provide our users with a way to manipulate a three-dimensional structure, but how can we do that without confusing our users? Neither Bush nor Engelbart discuss this issue because in their articles the structure is never sketched as having more than two dimensions. Looking for some structural tools for manipulating associative structures we quickly found the Topic Maps technology[6], ISO/IEC standard 13250:2000, partly because it is the technology used for storing the structure of the Polyscopic Navigator mentioned in the previous section. In Topic Maps we are able to connect topics of any kind through associations, and as this is similar to what Vannevar Bush wrote about they are well suited for the kind of structures he described. Rolf Guescini discusses using Topic Maps rather than RDF and Web Ontology Language (OWL) for his TMemex[11], and uses Topic Maps in the prototype appplication that implements the paths like the Memex. Two other tools that display the contents of a topic map are Ontopia's Vizigator[31] and TM View[40]. Vizigator has a 2D approach, while TM View uses 3D. What is appropriate for a collaborative web tool?

Let us start by looking at this using just two dimensions, and our example is the simple structure found in fig. 4 on the next page. It shows a 1:3-relationship between four elements, and element A is displayed as a level above elements B, C and D. In Polyscopy it would signify that the lower elements, B, C and D, have a higher level of detail than element A as they

[6]For a quick introduction to Topic Maps, see "The TAO of Topic Maps"[34].

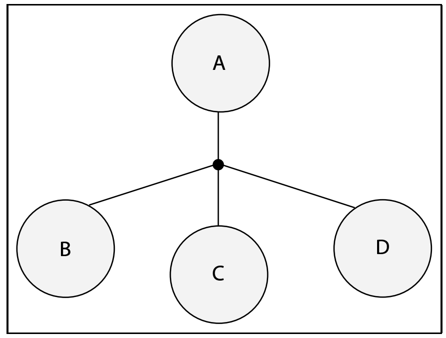

Figure 4: Vertical abstraction example viewed from the side

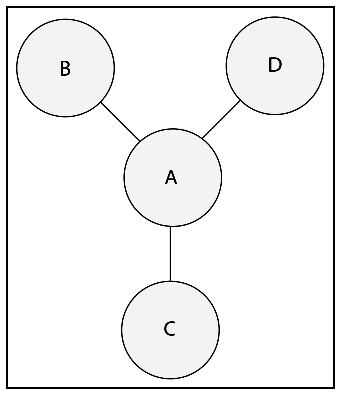

Figure 5: Vertical abstraction example in fig. 4 viewed from above

are one step down. It might also mean that element A is the overview, a document describing the bigger picture you get when combining the other elements. This approach does work, but look at what happens if we change our perspective so that the structure is viewed from above, as seen in fig. 5.

The structure is the same in both fig. 4 and fig. 5, but when we look at it from above its meaning changes drastically. The three elements B, C and D are still related to element A, but rather than having a central connection point they seem to have one connection each. It is also impossible to see what level each element is on, at first glance it seems they're all on the same level. Our conclusion is that displaying just two dimensions is a possibility, but having to switch between views like this to figure out how the elements are related is undesirable.

Another possibility is to have a three-dimensional application, and as we saw TM View is a 3D application for displaying the contents of topic maps. A drawback with their solution is that it uses a proprietary plugin. We believe it could also be possible to do this with a Java applet, for instance there are many examples of 3D visualizations to be found. Both are interesting possibilities as the 3D-approach could allow the user to rotate and move the structure around as they see fit.

In our case we believe there might be a better alternative by choosing the middle ground, what is dubbed 2.5D. The space we move around in is still three-dimensional, but we use only two dimensions to display it and allow the user for example a slider to move in the third dimension. Examples of this kind of navigations are aplenty, one is found in "Experiential design of shared information spaces"[14] and reported that the user's successfully navigated the three-dimensional space using the 2.5D approach. A similar approach that we'll come back to is the controls of map displays as seen by Google Maps and MapQuest.

We find the 2.5D approach appealing because as we saw with the example in fig. 3 on page 26 we believe creating and moving around on the levels should happen by connecting a complete structure to the overview. This means a view from above can be used since it will let us correctly lay out all elements on a level, and movement in the third dimension will only be between levels. Although 3D could be a nice feature it is not a must-have in the first version.

4 The Solution

On top of the foundation laid out in section 2 we started sketching out the rest of the structure in section 3. It is now time to move on to the specifics, see how the ideas of Bush and Engelbart can come together on the Web in a way that adheres to the principles of Polyscopy, while also utilizing the patterns described by Tim O'Reilly.

4.1 Main Components

For the web solution we have in mind we see four main components:

1. A user interface for working with symbolic structures
2. A way to connect the user interface to existing web pages
3. A search engine enabling users to search through the structure
4. Various ways of syndicating content from the system

We have chosen to leave the search engine out of the discussion. Although an important component, semantic search is not the main focus of this thesis and we also see it as dependant upon the choice of technology used for the system. There is much research being done on semantic search, and we believe we could better utilize that knowledge at a later stage should the system be created.

Remembering O'Reilly's patterns form section 2.5 two of the components are rather unsurprising, we see it as important that the system facilitates syndication of content. One reason is that the system should be more useful when the content is available in a myriad of formats, allowing people to use the content in new ways. Both graphical and textual representations should be useful, leaving it up to the one reusing the content to decide what is appropriate for their application. The second reason is that we think the added value can motivate users to contribute, which in combination with an interface that is not restricted to just our own system can increase the number of users and subsequently the amount of information created. Having this idea of user motivation in mind at an early stage is probably good, rather than create the whole system first and then figure out how to attract users to it.

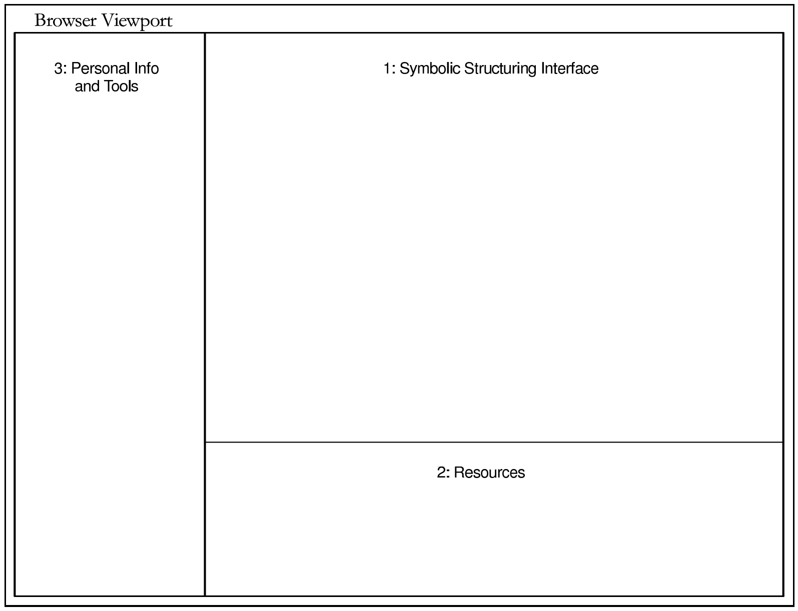

Figure 6: Sketch of the general user interface setup

4.1.1 User Interface

The basic idea of the major components of the user interface is shown in fig. 6, and as we can see it splits the browser viewport into three major pieces.

The main part of the interface is to be spent on the symbolic structuring interface which we will cover in more detail in the following sections. Below the symbolic map is an area to be used for connections to resources. This idea is similar to the Topic Maps sketches found in “The TAO of Topic Maps”[34] and it is kept this way because we see a need for separation between the symbolic structure and the actual resources it links to. Whether our users will be able to understand this difference, meaning if this solution actually works, needs to be determined through user testing.

The last part of the user interface is planned to be used for user information and various menu items, e.g. you will have a user profile with some statistics, menu items for getting a list of your friends, news, logging out, etc. Exactly what is needed and useful for this part of the interface is also something that could be determined in the design stage and through user testing. As we have not created a functional prototype this has not been completed. Instead, let us examine the symbolic structure and some key areas related to polyscopic structuring.

4.2 Symbolic Structuring on the Web

The ideas for the user interface brings together the ideas brought forward in 3.3.1 and the 2.5D concepts from 3.3.2. A closer look at the symbolic structuring and resource parts of the interface is shown in fig. 7 on the next page. We will start our description of the user interface as seen from the individual user's perspective, and then in section 4.4 cover the collaborative ideas found in the system.

We believe that reusing the concepts of a circle to symbolize a concept/topic and the line between the to symbolize the associative connection will make the user interface sufficiently intuitive. This representation is often seen in these kinds of applications, for instance it is also used in Ontopia's Vizigator that we mentioned earlier. We have added the controls for navigating the structure in a similar fashion to what is used in map application online, as discussed earlier. Movement in the two dimensions of the displayed structure is available through the N/S/E/W-controls, and movement in the third dimension, the level of detail, is available through the zoom slider. On the right hand side of the interface is a "toolbox" where we plan to place buttons for creating new concepts (circles), connections (lines) and joining these together. Thus mapping out the two-dimensional structure we have seen in our examples.

Additionally the users must be able to connect the concepts with resources like webpages that describe these concepts. At first glance it might seem strange that we have this distinction between concepts and the resources, and that we instead should allow the circle itself to be the reference to the resource(s). We believe that it will be beneficiary to keep the distinction between the resources and the symbolic structure above it, and that deciding whether to have circles that are a description or title of a resource or simply reference the resources directly should be left to our users. Either way will allow a concept to have many resources, and a resource to have many concepts, the main difference is only how you link them.

4.3 Creating Levels

Remembering our discussion in section 3.3.1 the challenge now is how add the third dimension to our structure. Our suggested solution is to have the users create a group, again through a button in the toolbox, and then connect that group to a higher level element. An illustration of this concept is seen in fig. 8 on page 34, where the state on the left is the "before" state and the one on the right is the "after" state. The process is easy to describe:

1. Create the new group

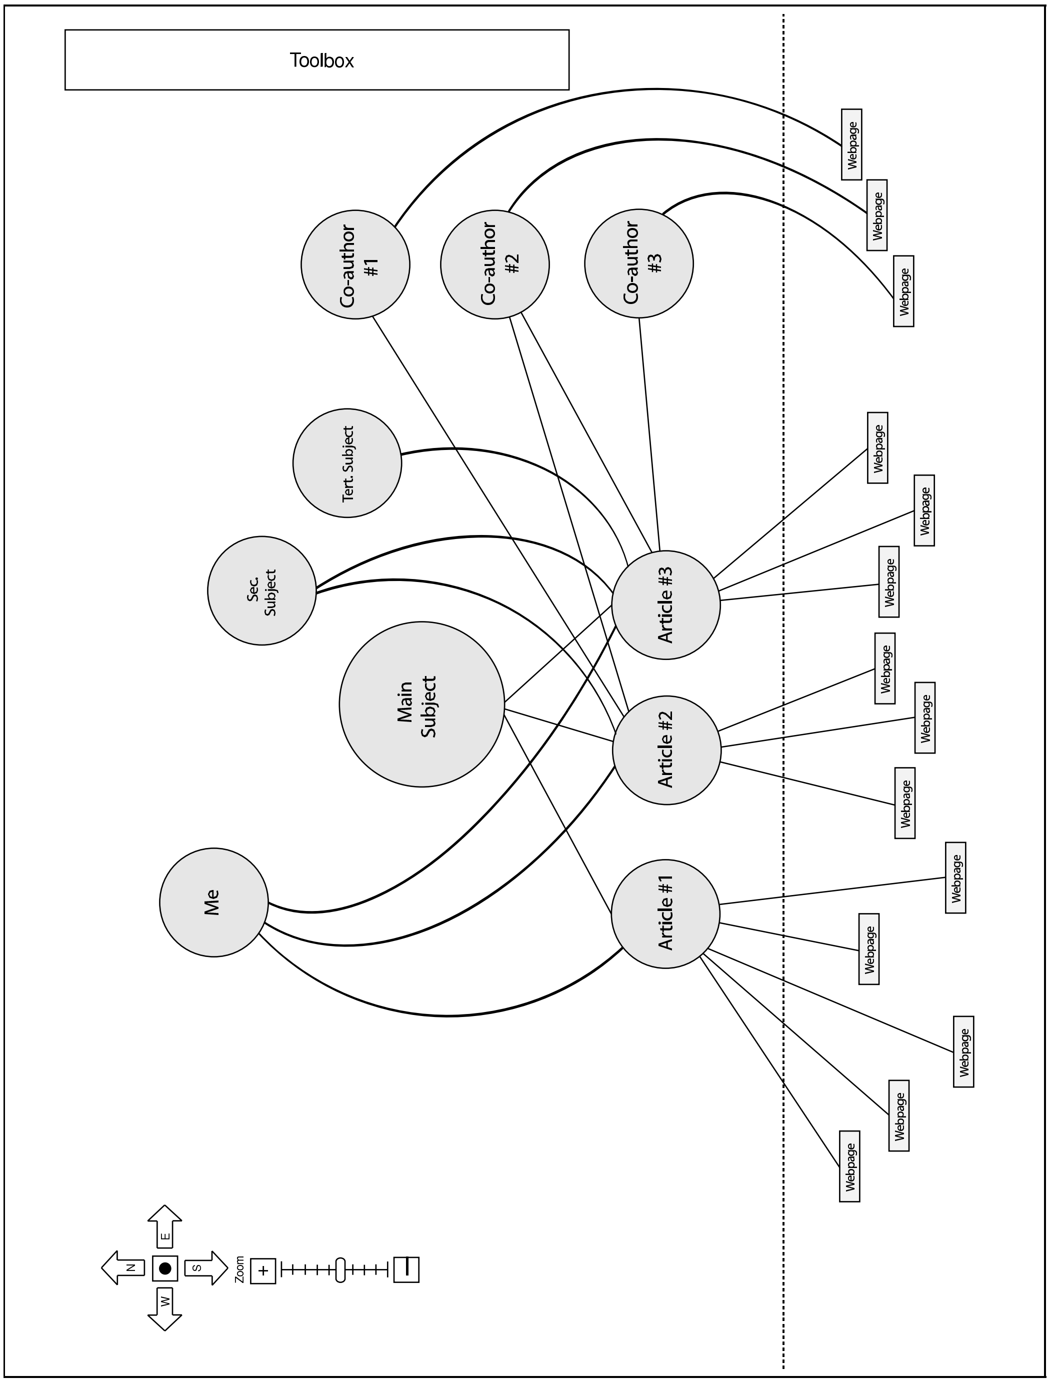

Figure 7: Sketch of the general user interface for symbolic structuring

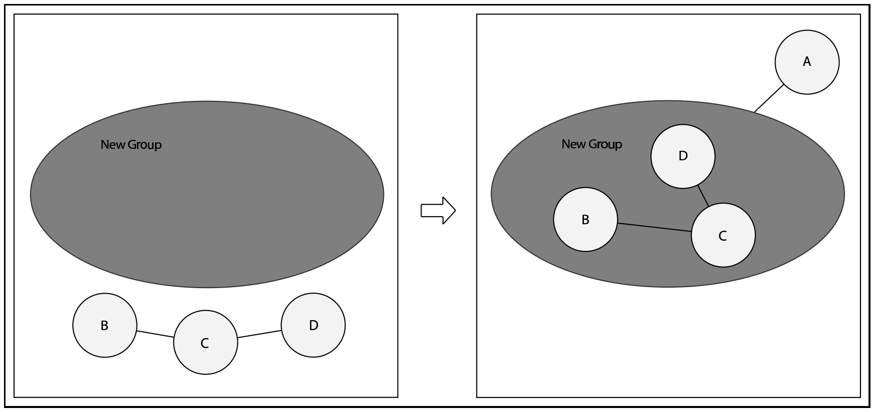

Figure 8: Illustration of the before and after states of creating a group

2. Add items to the group, drag'n'drop is a possibility
3. Associate the group with the overlying concept

Having the lower level structure, moving it into the group and connecting it to a higher level element is the bottom-up way of creating the three-dimensional structure. A top-down approach is also possible by first creating the top level element (element A in fig. 8), then creating the group and the connection to the top level element, and finally by adding the lower level structure to the group as needed.

In our example illustration the contents of the group are shown but because it changes the structure from having three distinct dimensions to having two dimensions and groups our plan is to not show the contents in the application. For larger and more complex structures it is easier to argue that they shouldn't be shown as space inside the group restricts the view, and secondly that it could be confusing to the user.

4.4 Collaboration

So far we have only inspected the interface from the perspective of the individual user. It is time to let the rest of the world on the field. First of all we see that the system needs a way to easily share information with others. An illustration of how we envision the process of doing that is seen in fig. 9 on page 36. The main idea behind the process is that changing the visibility is similar to choosing which side of a fence something is on.

In the illustration in fig. 9 on page 36 we see four elements similar to those from the create group example in fig. 8, named A, B, C and D. The three latter elements and their relationship are already visible to everyone, while element A is moved to a middle ground. A group of co-authors have been created and the user allows the people in that group to see element A by moving the element into the space where the group is.

Implicit in the illustration in fig. 9 on the next page is the idea that groups can also consist of people, and they might even be users of the system. Creating such a user group is done just like creating any other group of elements as we saw in fig. 8 on the facing page. This means that the concept of a friend on social sites that we discussed in section 3.1 can be recreated by a group, perhaps named “My friends”, where you move the elements that are your friends into. One neat feature of doing it this way is that it is very easy to move people around, either in and out of a group or between them. Creating groups of groups should of course also be possible.

In addition to being able to change what other people see, we also believe it should be similarly easy to change what other people can work with you on, and believe it can be done in just the same fashion, and that there should also be a way to change both permissions at the same time.

With the possibility of sharing and working on the structures there comes a need for communication. We believe in allowing many ways of communication, and therefore see at least three different possibilities:

1. Adding notes to any resource, these might be just your own ideas about something, similar to sticking a Post-It Note on something.
2. Shared notes, where those who can see and/or edit an element can do quick discussion of minor issues.
3. Discussion pages, similarly to those in Wikipedia, that are a page dedicated to communication about a specific element in the system, and suitable for more ellaborate discussion.

An illustration of how a shared note could appear in the system is seen in fig. 10 on page 37.

Although we suggest to recreate some useful features from Wikipedia it is important to note that the system is not designed to be geared towards creating an encyclopedia. Wikipedia has a principle of Neutral Point of View (NPOV) and through communication editors of an entry seek to reach a common agreement on what an unbiased presentation of a subject is. The users of the system we described might choose to work in a similar way, or they can choose not to. Either way should be welcome, and rather than moving back and forth between revertions of an entry they can agree to disagree and leave their maps they way there were. They could even end up showing this situation to the users, an illustration of a disagreement is shown in fig. 11 on page 38.

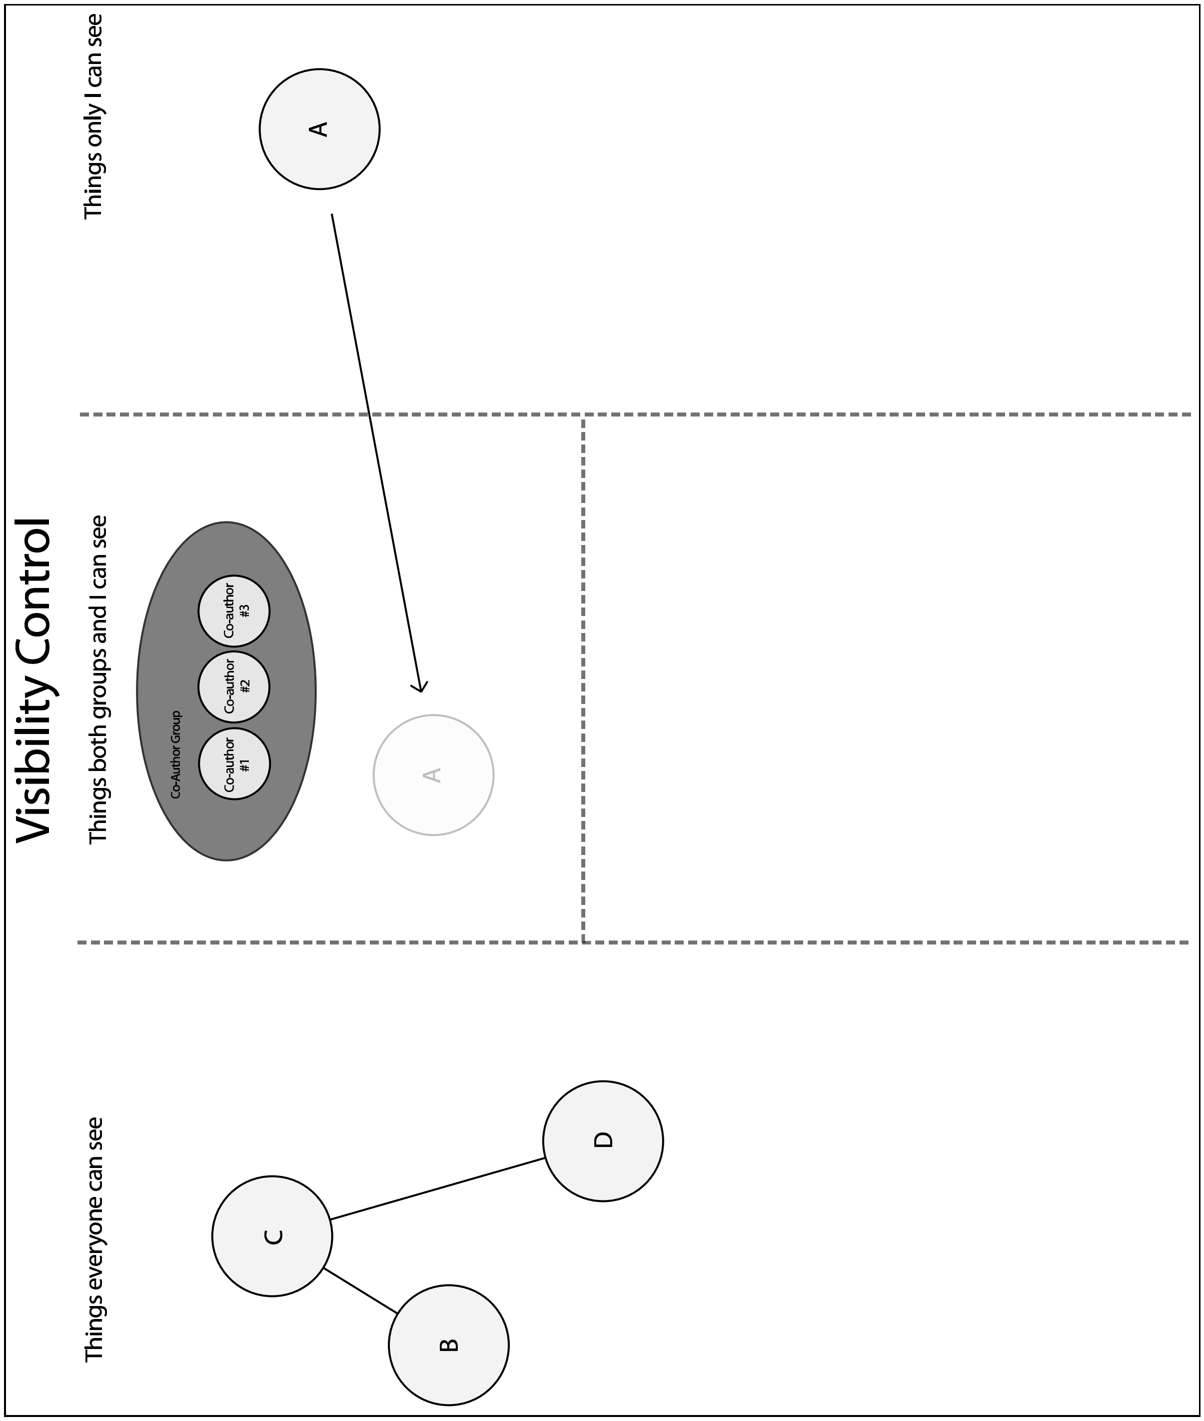

Figure 9: Illustration of how to change visibility of an item

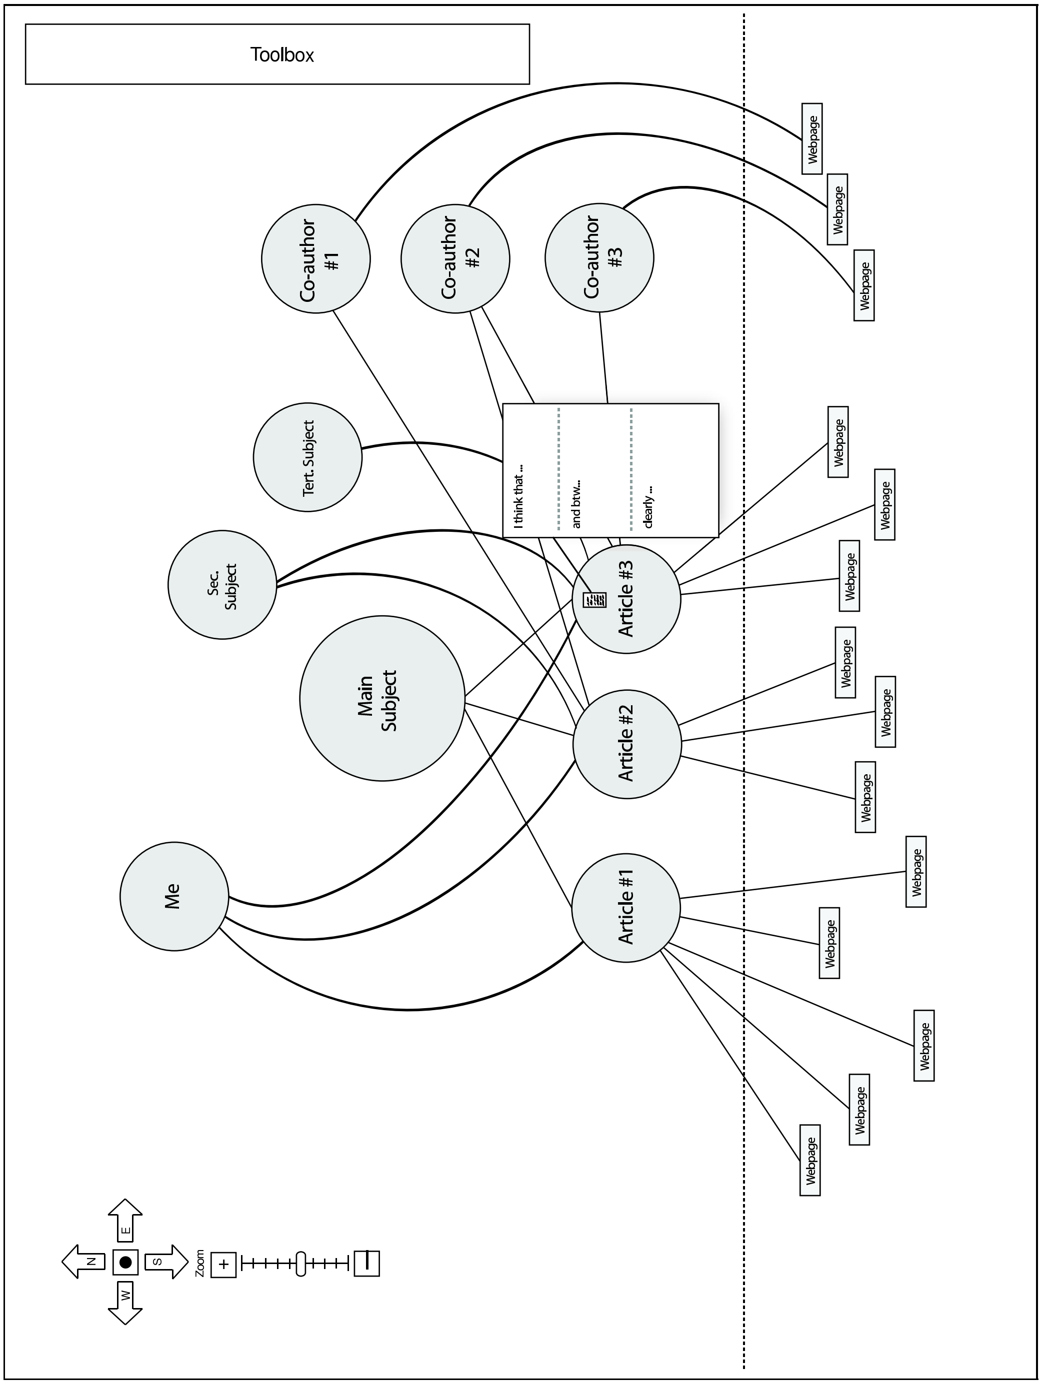

Figure 10: Illustration of how a discussion about an element can be displayed

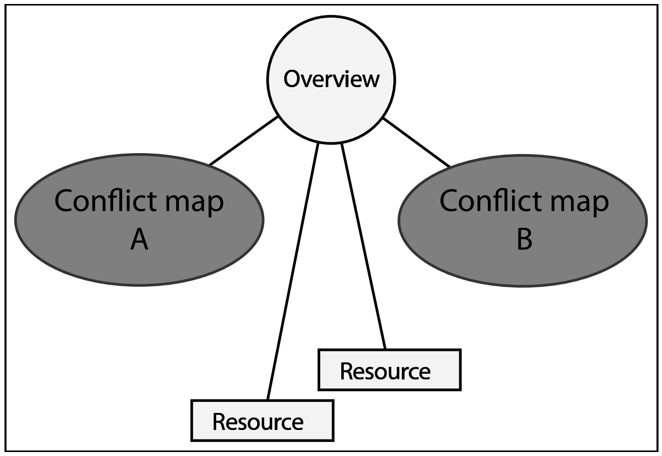

Figure 11: Illustration of displaying disagreement through two connected maps and resources

In fig. 11 the two different opinions have created two groups where their respective structures are put. An overview element is then created and attached to the groups, and it in turn refers to a webpage with a textual presentation of the issue at hand, perhaps also with the possibility of continued discussion. We don't see everyone agreeing with each other at all times as a possible outcome, and instead we try to design a system that allows people to keep on disagreeing, but at the same time also providing them with a possibility of mapping out the disagreement in a structure.

4.5 Paths

One of the most interesting features of Vannevar Bush's Memex is perhaps the associative paths. An example implementation is Walden's paths[36, 9, 14] and Rolf Guescini also implemented path features for the TMemex application in his master's thesis[11].

Our vision of how the application can display a path to the user is seen in fig. 12 on the facing page. For a path to be possible the dimension of time needs to be added, and in our case we use a timeline, similar to those found in video editing software. In the center of the display is the currently selected element, "Main Subject", shown together with its relation to other elements. The group of four elements connected to the left of the main subject are the earlier elements in the path, while the group of six elements to the right are the future elements. This particular path therefore consists of eleven elements all connected in a straight line, and the timeline moves from left to right.

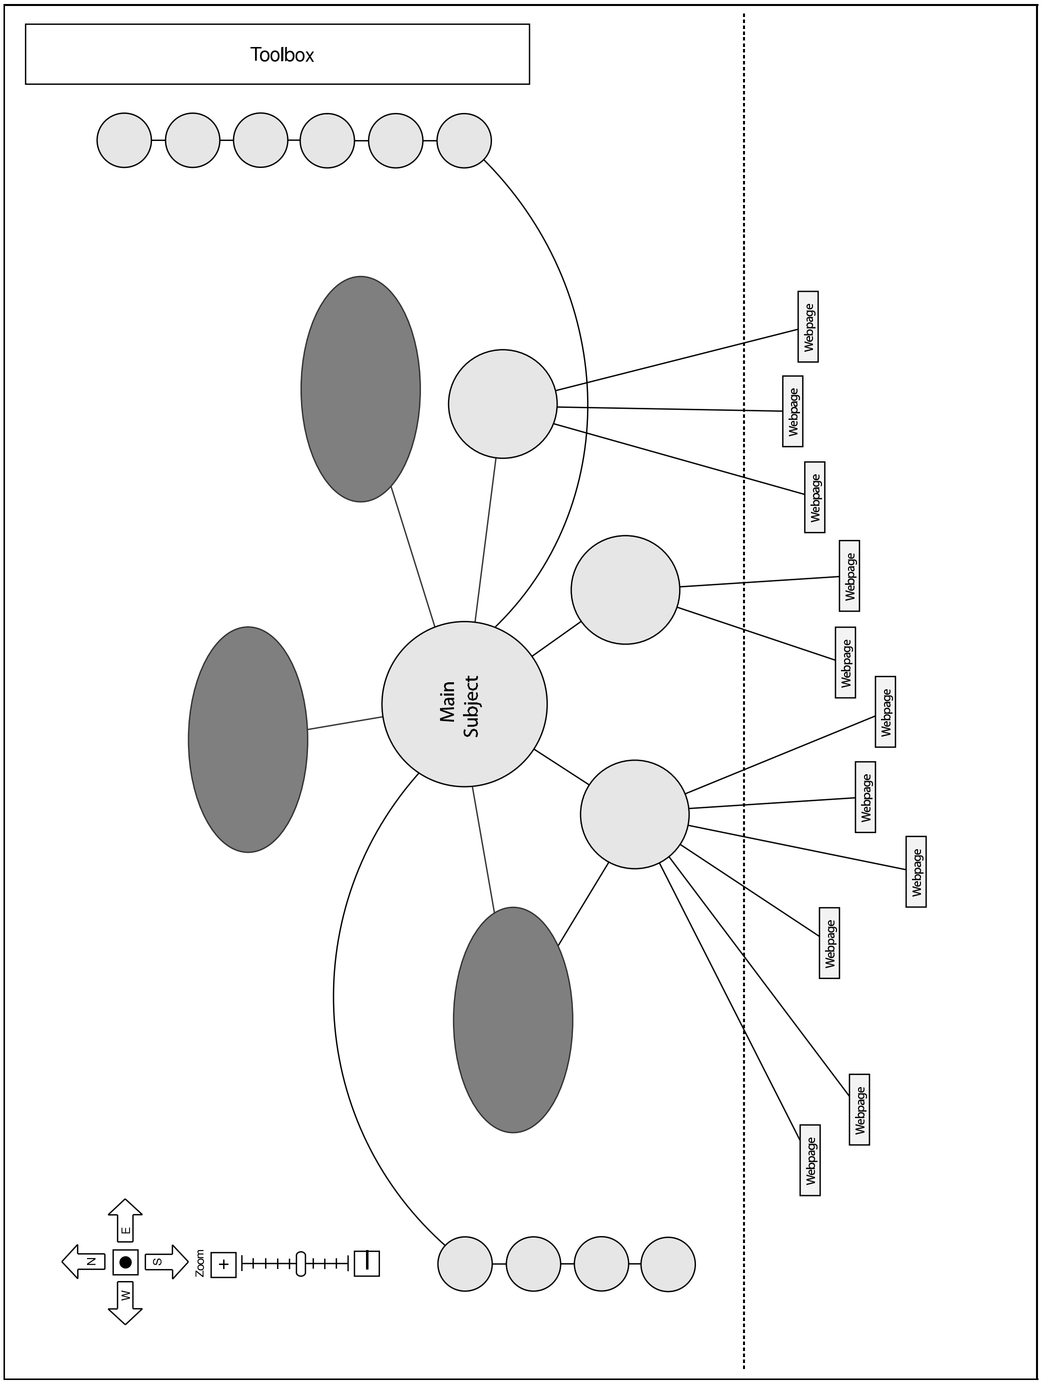

Figure 12: Illustration of walking a path

One of the really interesting collaborative features described by Vannevar Bush in “As We May Think” is the exchange of the associative paths. This means there needs to be a way of storing them and allowing others to get access to them. We also see the need for ways of extending them, meaning users should be able to continue down a different connection in the structure, which would split the path. They could then store and annotate it so that others can learn from what they have found as they went along.

We also see a possibility for a more traditional approach to information gathering, using a browser plug-in to store the navigation history. That history is also a path, or a set of paths, and this can then be uploaded to the site where the user can annotate and share the whole history or parts of it. If we add the ability to search for paths and a way to record, store and index audio and video we end up with something quite close to the Memex Bush envisioned.

5 The Future

Having described a system for allowing people to create and collaborate on symbolic structures the question is where do we go next? First on the list of answers is the obvious one, build it and see if it works. In this project too much time was spent on research and figuring out a way that things could work, rather than make a prototype testbed. We believe that the system can be built using currently available web technology, both Java and Macromedia Flash should have the capability to create the user interface. A cutting edge approach would be to attempt to create it using JavaScript[7] and Scalable Vector Graphics (SVG), which might be ahead of its time. On the server side we lean towards the Topic Maps technology because it comes with the associative structures, and it also has other mechanisms for handling challenges like internationalization and grouping. The former can be tackled through Topic Maps' ability to have multiple names for everything while the latter can be solved with scopes, which is Topic Maps' mechanism for filtering out irrelevant parts of the structure.

In addition to the obvious answer we have also found two other areas of particular interest.

5.1 Identity management

From the illustration of altering permissions that we showed in section 4.4 it follows that the symbolic structure can also include references to actual people, meaning that the element that has my name on it is a reference to the physically existing person that is me. This creates a lot of interesting possibilities for mapping out my own world in a way that I understand, which is part of the motivation behind this thesis, allowing the individual user to do just that. It also creates a lot of interesting challenges, like how to verify that I am who I claim to be, and can we make it possible to handle all information that refers to me without it becoming unmangeable? There would also need to be a way to take care of multiple representations of me, meaning that if someone else creates an element that is supposed to refer to me I can seize control of it and merge it with the actual me. Unless we find that it is better to let go of the control and allow more than one me to exist.

[7] The standard is called ECMAScript (ECMA-262), but the popularly used name is JavaScript.

5.2 Dynamic Units of Information

Going back through history can reveal interesting things about how our ways of storing and accessing knowledge have changed. Scrolls were a linear medium where referencing specific parts is was, while books make it easier to refer to specific places as they have pages. With movable type came the ability to mass produce books and as copies were identical to each other it allowed referring to particular passages a breeze.

Today we have powerful information technology available, is it time we made a similar transition when it comes to how knowledge is created and accessed? We still create a lot of information and store it in large units, for instance in books or documents like this thesis. Can we create a way of accessing and transferring knowledge that is based on smaller units? Similar to how David Weinberger describes what happened to CDs in "Everything is miscellaneous"[47][8]:

> "As soon as music went digital, we learned that the natural unit of music is the track."

What is the natural unit of knowledge? Is it a book, a document, a page, a paragraph, or perhaps just a sentence?

One of the ideas found in Polyscopy is that we should create information that allows the user to choose between several difference scopes, or points of view. A radical way to do that would be to have documents that display differently depending on what you are seeking to find there, or maybe also depending on where you come from. Have information that changes with the context it's presented in, and by that we don't mean simply altering the presentation slightly, we mean changing the content. We could also allow the user to choose between different scopes and change the content on the fly, letting them see how the message changes with the point of view. It would mean we move away from having the document as the unit and instead focus on the subject it should present, or making a paradigm shift from document-centric computing to subject-centric computing[30].

Going back in history can also mean you find some radical ideas for information technology. In "Glut" by Alex Wright[55] we find a chapter called "The Web That Wasn't"[9]. we can discover Paul Otlet, a Belgian who already in 1934 sketched out a plan of a global network what he called "electric telescopes". His vision was interlinked devices enabling searching and browsing

[8] See also Weinberger's Google Tech Talk on "Everything is misc." on YouTube[22].

[9] Wright also did a Google Tech Talk on "The Web That Wasn't" available on YouTube[23], and for a short introduction to Paul Otlet's work see "The web that time forgot"[56]

of millions of documents, images, audio and video. The difference between his vision and our World Wide Web is that his sketches show technology not so distance from what he had available.

Otlet's most interesting work is perhaps not the vision of a network, he also attempted to create an index of all the knowledge in the world. It was known as the Mundaneum, and its service to the world was a search engine driven by humans, handling more than 1500 queries a year. That is not a large number by today's standard, but not only do have we substituted humans with technology, we also often use a different service. Where a Web search might provide you with ten pages you have to check out yourself, the Mundaneum was slightly more of an oracle. Otlet had created a way of indexing information that not only had cross-references, it also added meaning to the references, you could find out whether a reference was used for support or disagreement.

What Otlet had in mind was to liberate knowledge from the boundaries of the book. The question for us is then, can we liberate knowledge from the boundaries of the Web?

6 Acknowledgements

I would like to thank my family and friends for their love and support through both ups and downs as I worked my way towards completing this thesis.

I would like to thank Dino Karabeg, my advisor, for creating the Polyscopic methodology I have explored during my work, and for his ideas and help with the thesis.

I would like to thank all my instructors both at the University of Oslo and the University of Minnesota for allowing me to benefit from their knowledge and experience.

7 Appendices

List of Figures

References

[1] Tim Berners-Lee and Mark Fischetti. *Weaving the Web: The Original Design and Ultimate Destiny of the World Wide Web*. Collins, 11 2000.

[2] Z. Bokhari. Standard & poor's industry surveys, computers: Software. `http://www.netadvantage.standardandpoors.com/NASApp/NetAdvantage/showIndustrySurvey.do?code=cos`, October 2006.

[3] Al Borchers, John Herlocker, Joseph Konstan, and John Riedl. Ganging up on information overload. *Computer*, 31(4):106–108, April 1998.

[4] Vannevar Bush. As we may think. `http://www.theatlantic.com/doc/194507/bush`, July 1945.
Retrieved on July 5 2008.

[5] Alan Cooper. *About Face - The Essentials of User Interface Design*. IDG Books Worldwide, Inc., 1995.

[6] Cory Doctorow. Transcript of 'folksonomy, or how i learned to stop worrying and love the mess'. `http://craphound.com/etech2005-folksonomy.txt`, March 2005.
Retrieved on July 15 2008.

[7] Barbara Ehrenreich. *Bait and Switch - The (Futile) Pursuit of the American Dream*, chapter 1, pages 29–30. Metropolitan Books, 2005.

[8] Douglas Engelbart. Augmenting human intellect. `http://www.bootstrap.org/augdocs/friedewald030402/augmentinghumanintellect/ahi62index.html`, October 1962.

[9] Richard Furuta, III Frank M. Shipman, Catherine C. Marshall, Donald Brenner, and Hao wei Hsieh. Hypertext paths and the world-wide web: experiences with walden's paths. In *HYPERTEXT '97: Proceedings of the eighth ACM conference on Hypertext*, pages 167–176, New York, NY, USA, 1997. ACM.

[10] David Goldberg, David Nichols, Brian M. Oki, and Douglas Terry. Using collaborative filtering to weave an information tapestry. *Commun. ACM*, 35(12):61–70, 1992.

[11] Rolf Guescini. Getting back on the trail, combating information overload with topic maps. Master's thesis, University of Oslo, May 2006.

[12] Rolf Guescini, Dino Karabeg, and Tommy Nordeng. A case for *Polyscopic* structuring of information. *Lecture Notes in Computer Science*, 3873/2006:125–138, February 2006.

[13] Tim Harford. *The Logic of Life: The Rational Economics of an Irrational World*. Random House, January 2008.

[14] Kristina Höök, David Benyon, Alan J. Munro, Dan Diaper, and Colston Sanger, editors. *Designing Information Spaces: The Social Navigation Approach*. Springer-Verlag, London, UK, 2003.

[15] Giovanni Iachello, Ian Smith, Sunny Consolvo, Mike Chen, and Gregory D. Abowd. Developing privacy guidelines for social location disclosure applications and services. In *SOUPS '05: Proceedings of the 2005 symposium on Usable privacy and security*, pages 65–76, New York, NY, USA, 2005. ACM. Computing.

[16] 280 North Inc. 280 slides. http://280slides.com/.
Retrieved on July 4 2008.

[17] Alexa Internet Inc. Alexa top 500 sites. http://www.alexa.com/site/ds/top_sites?ts_mode=global&lang=none, 06 2008.
Retrieved on June 18 2007.

[18] CondeNet Inc. reddit.com. http://www.reddit.com/.
Retrieved on June 20 2008.

[19] Digg Inc. Digg. http://www.digg.com/.
Retrieved on June 20 2008.

[20] Digg Inc. Digg the blog » blog archive » digg friends. http://blog.digg.com/?p=41, 09 2006.
Retrieved on June 23 2008.

[21] Google Inc. Google docs. http://docs.google.com/.
Retrieved on July 4 2008.

[22] Google Inc. Everything is miscellaneous (google tech talks). http://www.youtube.com/watch?v=x3wOhXsjPYM, October 2007.
Retrieved on June 18 2008.

[23] Google Inc. The web that wasn't (google tech talks). http://www.youtube.com/watch?v=72nfrhXroo8, October 2007.
Retrieved on June 18 2008.

[24] Erlend Ofte Johansen and Johannes W. Berg. Wikipedianerne. Newspaper article in Dagens Næringsliv, July 26/27 2008.

[25] Netcraft Ltd. Web server survey archives. http://news.netcraft.com/archives/web_server_survey.html, June 2008.
Retrieved on July 29 2008.

[26] CNN Money. Yahoo! buys geocities. http://money.cnn.com/1999/01/28/technology/yahoo_a/, January 28 1999.
Retrieved on July 29 2008.

[27] Mark R. Nelson. We have the information you want, but getting it will cost you!: held hostage by information overload. *Crossroads*, 1(1):11–15, 1994.

[28] CNET News.com. Google buys remaining deja.com business. http://news.cnet.com/2100-1023-252449.html, February 12 2001.
Retrieved on July 30 2008.

[29] Jakob Nielsen. Participation inequality: Encouraging more users to contribute. http://www.useit.com/alertbox/participation_inequality.html, October 2006.
Retrieved on July 20 2008.

[30] Ontopedia. Subject-centric computing. http://psi.ontopedia.net/Subject-centric_computing.
Retrieved on August 1 2008.

[31] Ontopia. Ontopia solutions – vizigator. http://www.ontopia.net/solutions/vizigator.html.
Retrieved on July 29 2008.

[32] Tim O'Reilly. What is web 2.0. http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html, 09 2005.

[33] Paul. Housingmaps. http://www.housingmaps.com/.
Retrieved on July 4 2008.

[34] Steve Pepper. The tao of topic maps. http://www.ontopia.net/topicmaps/materials/tao.html, June 2000.
Retrieved on July 29 2008.

[35] R. Priedhorsky, J. Chen, S.T.K. Lam, K. Panciera, L. Terveen, and J. Riedl. Creating, destroying, and restoring value in wikipedia. *Proceedings of the 2007 international ACM conference on Conference on supporting group work*, pages 259–268, 2007.

[36] Walden's Paths Project. Walden's paths project homepage. http://www.csdl.tamu.edu/walden/.
Retrieved on July 31 2008.

[37] Paul Resnick, Neophytos Iacovou, Mitesh Suchak, Peter Bergstrom, and John Riedl. Grouplens: an open architecture for collaborative filtering of netnews. In *CSCW '94: Proceedings of the 1994 ACM conference on Computer supported cooperative work*, pages 175–186, New York, NY, USA, 1994. ACM.

[38] Shilad Sen, F. Maxwell Harper, Adam LaPitz, and John Riedl. The quest for quality tags. In *GROUP '07: Proceedings of the 2007 international ACM conference on Supporting group work*, pages 361–370, New York, NY, USA, 2007. ACM.

[39] Shilad Sen, Shyong K. Lam, Al Mamunur Rashid, Dan Cosley, Dan Frankowski, Jeremy Osterhouse, F. Maxwell Harper, and John Riedl. tagging, communities, vocabulary, evolution. In *CSCW '06: Proceedings of the 2006 20th anniversary conference on Computer supported cooperative work*, pages 181–190, New York, NY, USA, 2006. ACM.

[40] Spatialknowledge.com. ::spatialknowledge:: a xtm compliant topicmap viewer. `http://www.spatialknowledge.com/projects/topicmaps/`. Retrived on July 29 2008.

[41] James Surowiecki. *The Wisdom of Crowds*. Anchor, August 2005.

[42] Telegraph.co.uk. Youtube lawsuit threatens internet freedom, claims google. `http://www.telegraph.co.uk/connected/main.jhtml?xml=/connected/2008/05/27/dlgoog127.xml`, May 27 2008.
Retrieved on July 16 2008.

[43] F.B. Viégas, M. Wattenberg, and K. Dave. Studying cooperation and conflict between authors with history flow visualizations. *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 575–582, 2004.

[44] Jesse Vig. geogreeting. `http://geogreeting.com/`.
Retrieved on July 4 2008.

[45] W3C. World wide web consortium. `http://www.w3.org/`.
Retrieved on July 4 2008.

[46] Morten Warncke-Wang. Polyscopic document navigator. `http://www.ifi.uio.no/info-design/projects/2005/nettrom/web-site/`, 05 2005.
Retrieved on June 26 2008.

[47] David Weinberger. *Everything is miscellaneous*, chapter 1, page 9. Times Books, 2007.

[48] Wikipedia. Deja news. `http://en.wikipedia.org/wiki/Deja_News`.
Retrieved on July 30 2008.

[49] Wikipedia. Digg – wikipedia. `http://en.wikipedia.org/wiki/Digg`.
Retrieved on July 29 2008.

[50] Wikipedia. Wikipedia:edit war. http://en.wikipedia.org/wiki/Wikipedia:Edit_war.
Retrieved on July 30 2008.

[51] Wikipedia. Wikipedia:errors in the encyclopædia britannica that have been corrected in wikipedia. http://en.wikipedia.org/wiki/Errors_in_the_Encyclopedia_Britannica_that_have_been_corrected_in_Wikipedia.
Retrieved on July 29 2008.

[52] Wikipedia. Wikipedia:external peer review. http://en.wikipedia.org/wiki/Wikipedia:External_peer_review#Nature_.28December_2005.29.
Retrieved on July 29 2008.

[53] Wikipedia. Wikipedia:external peer review/nature december 2005/errors. http://en.wikipedia.org/wiki/Wikipedia:External_peer_review/Nature_December_2005/Errors.
Retrieved on July 29 2008.

[54] Wikipedia. Wikipedia:list of policies. http://en.wikipedia.org/wiki/Wikipedia:List_of_policies.
Retrieved on July 30 2008.

[55] Alex Wright. *Glut: Mastering Information Through The Ages*. Joseph Henry Press, June 2007.

[56] Alex Wright. The web that time forgot. http://www.iht.com/articles/2008/06/17/healthscience/17mund.php, June 17 2008.
Retrieved on August 1 2008.