Research at the table

Bram Vandeputte and Erik Duval

Katholieke Universiteit Leuven, Department of Computer Science
Celestijnenlaan 200a, 3001 Leuven, Belgium
{bram.vandeputte,erik.duval}@cs.kuleuven.be

Abstract In this paper we describe how we want to take advantage of the rapid developments in technology to assist researchers in doing research. More specifically in exploring the publication space. For this purpose we have designed and developed a prototype application to take advantage of large displays with multi touch enabled input. We describe the current state, the next steps and how we will to evaluate it. To conclude we give an outlook on further possibilities and challenges that lay ahead.

Key words: research2.0, information visualization, multi touch, large display, research.fm

1 Introduction

How great would it be to integrate the process of exploring publications, finding them and reading them in an almost seamless way ? This sort of idea was already described in *As we may think* in 1945 by Vannevar Bush [2]. The Memex was described as the perfect desk of a researcher, having all the knowledge of the world readily available. At that time personal computers were not even invented, but since then technology has advanced tremendously and become very common. Using current state of the art technologies, we want to find out how we can ease the process of exploring publications. This process is an important part of a researcher's job, as he wants to know what is going on in his field of research.

To be able to get this kind of understanding, Russell et al [10] have pointed out that it is imperative that the right representation is found for exploring a network of (publication) data.

The idea of visualizing publication networks has been inspired by the work of Klerkx et al, where they explore learning object repositories [8] and social bookmarks [7] in a visual manner.

In this paper we first introduce and describe the problem. We then motivate our hardware platform, describe the origin of the data and we explain the detailed workings of the application. In the next section we compare our work with existing studies. Then we describe how we evaluate this and finally we propose the next steps to be taken. To conclude we summarize our findings and discuss further possibilities.

2 Problem statement

An important part of a researcher's job is reading scientific papers. This ensures that the researcher is up to speed of what is going on in his research field. It is also a prerequisite for writing scientific papers, as handbooks such as the one from Robert A. Day [3] emphasizes.

There are three basic ways of dealing with scientific papers. There is active search, where you search for a particular paper or a 'good' paper on a specific topic you have in mind. There are dozens of websites that serve this purpose really well, such as Google Scholar[1], ISI web of Knowledge[2], DBLP[3] etc. There is also what we can call passive search, where you get alerted whenever new publication material is available. Google Scholar has recently added a feature where you can be alerted whenever something new comes up that matches certain keywords. Also many of the journal magazines let you subscribe to a list to send you the table of contents when a new issue is available. Finally one can focus on relations between papers and authors. There are existing tools where this is possible, but we think that there is not enough technical support available for exploring these networks.

To explain the problem we want to solve, we will briefly describe the use cases we want to tackle with this work. The use cases can be grouped into two categories. In the first category the use cases have a mainly *top down* approach, while the second category holds the use cases that typically need a *bottom up* approach.

2.1 World overview

Typically, in this use case a user would like to start with a complete overview of all nodes laid out in a graph. The user then wants to zoom in on parts of the graph that draw her attention. This can be used to find out patterns or clusters. In this case the user usually is already an expert in the field, trying to understand or improve his knowledge about the field.

2.2 Explore your neighbors

In this case you might want to start from a view with a focus on yourself, or the author or paper that you want to start from. Then you want to browse to nodes in your 'neighborhood', which are likely to be related and/or interesting. Here you can try to find answers to questions like : Where am I in the research publication space ? Who should I talk or connect to ?

[1] http://scholar.google.com
[2] http://apps.isiknowledge.com
[3] http://www.informatik.uni-trier.de/~ley/db/

3 The application

3.1 The hardware

The input modalities We chose for supporting a multi-touch setting, as we want to explore direct and multi touch capabilities. This to find out whether these relatively new input methods can help to make it easier for researchers to interact with the fairly complex graph like structures.

The display The application will entail a visualization of a deeply connected network containing up to hundreds (maybe thousands) of nodes. This property feeds the need for using a large display. These large displays, with increasingly higher resolutions, are also rapidly becoming cheaper and more common, which makes it easier to include them in our study and makes this study more relevant.

A problem that sometimes arises on multi touch input devices is when one touches the screen to give input, the finger or hand occludes information one wants to see at that moment. This can be solved in two ways, either we make the information appear next to the touch point, or we make the information bigger so it is less likely to be occluded. Both solutions can benefit from a larger display, as you have simply more space to put the information.

Studies by Forlines et al [5] and Kin et al [6] have already shown that on tabletop displays multi touch input has performance and spatial awareness advantages over the traditional mouse, which reinforces our choice of hardware. From a research perspective, we want to explore if and how a large screen estate can influence the possibilities of this kind of visualization.

3.2 The data

EC-TEL conference Our first scope was to visualize all the publications from all editions of one conference. We extracted metadata from papers and put them in a database. Unfortunately this extraction process is still very error prone and a lot of semi-manual cleaning up needed to be done. The approach took quite a bit of effort and is not very scalable.

To try and make access to these publication data easier, we propose an open architecture for exchanging these publication metadata. This architecture is currently being discussed and developed in the STELLAR project[4], with both suggestions for collecting these data using BuRST feeds[5] and a webservice API, called research.fm[6], to make them available for tools and widgets like the one we are describing in this paper.

[4] http://www.stellarnet.eu/
[5] http://stellarnet.eu/d/6/3/BuRST_format_adaption_discussion
[6] http://www.stellarnet.eu/d/6/3/KULDocumentation

3.3 The network and the visualization

The obvious relations to visualize are the paper-author relations, and also co-authorship. To build up this network, we want to have a self-organizing and self-declutering algorithm. We chose to use Traer Physics[7], an implementation of a simple particle system physics engine, which allows to combine a spring-graph algorithm with physical forces. This combination will take care of the organizing and declutering of the network, so we don't have to care about where to put the nodes. After experimenting with the parameters such as force, drag, mass of the particles, spring length and strength, ... We could see a clear network-like graph appearing when the network is stabilizing after a few seconds.

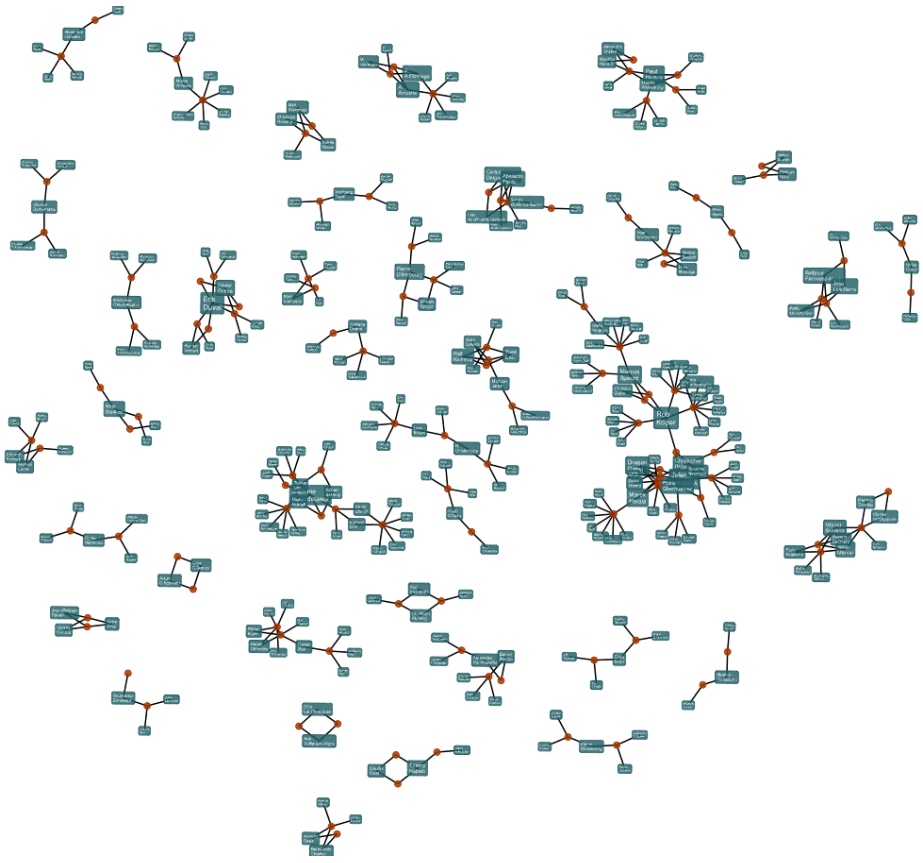

Figure 1: Overview of the whole publication network. The green nodes are authors, while the red ones are papers.

[7] http://www.cs.princeton.edu/∼traer/physics/

Figure 1 shows a screenshot of the visualization in the overview state. All nodes present in the network are shown. This state addresses the first use case we described in section 2.1. It can help a researcher to find out whether there is a lot of collaboration going on in this field, where the biggest clusters can be found or who the most active authors are.

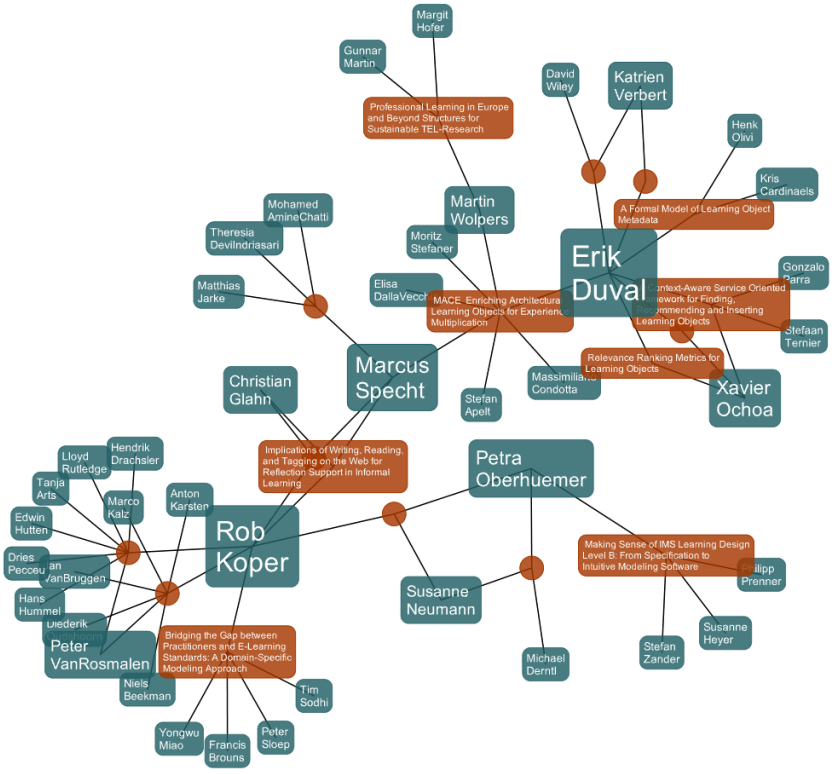

Figure 2: A detailed view of related authors. The green nodes are authors, the bigger they are, the more papers they have published. The red nodes are papers, where some of them have been expanded to show the title of the paper.

The second use case described in section 2.2 benefits from the view as shown in Figure 2. Here the visualization is zoomed in on a specific target. All the author names become clearly visible, so you can find an author very relevant for your work. One can also click on some paper nodes to get more information on the paper itself, so to find papers that are interesting, for example because they are closely related to your work. As you can see we are already experimenting with varying the node size of the author, based on his number of publications, to denote importance of this author.

4 Related Work

There are numerous other visualizations of publication data existing already. In this section we will highlight some visualizations that try to solve similar problems, and we will shortly describe how each of them differ from our approach.

4.1 Papercube

When this web application[8] first opens up, it immediately shows you a search box. This is useful when you are looking for something more specific, but it does not help when you want to explore the publication space and don't have a specific entry point in mind. There are quite some possibilities both in terms of relations and type of visualizations, so it can take a while for someone to get used to the interface and find what one actually wants. In our approach on the other hand, we want to make it easy for starting the exploration phase by directly showing the data. In this visualization the data is shown in a spring graph with a good lay-out. When you hover over a paper, the relations to other papers are highlighted, which is very helpful. One can also directly click through to the paper itself, so if you have found an interesting publication you can directly retrieve it online. Bergström et al [1] evaluated this application, and found that the users unanimously said that this kind of visualization can usefully augment existing digital libraries.

4.2 Ed-Media Relation Browser

The Ed-Media Relation Browser[9] is also an interactive, browser based, author visualization. In this approach they focus on one person and its direct relations, assisted with a strong filtering mechanism. The visualization only starts after you have entered a name. This emphasizes their focus on solving the problem of getting to know closely related authors. It does not allow one to study the field nor to discover the indirect relations between authors and papers. In our approach we try to solve this problem by allowing to zoom in on a specific person, but with a global navigation strategy so that the overview does not get lost. This visualization does not allow to rearrange the graph. To help the spatial memory we allow the user to organize the papers and authors however he likes. The authors, Ochoa et al [9], have also studied the complete publication space of a conference, but only with non interactive visualizations, where we allow to do so with a highly interactive visualization.

4.3 Microsoft Academia Search Visual Explorer

The Microsoft Academia Explorer[10] is similar to the Edmedia Relation Browser. Here you can drag the authors around to get a better view if something is not

[8] http://papercube.peterbergstrom.com
[9] http://ariadne.cs.kuleuven.be/edmedia/
[10] http://academic.research.microsoft.com/VisualExplorer.aspx

clear. Once you click on an other author, the graph keeps the link with the previous author but unfortunately all not directly related authors get thrown away. Thus also this visualization only displays direct relations. This application is also only targeted at visualizing authors. One can click through to see all the details of an author, but it is not possible to see the publication which make authors related. Our approach makes the transition from exploration to reading papers easier by bringing the papers visually in the network. If a paper draws attention, one can immediately retrieve more information from it.

5 How to evaluate ?

Due to the early stages of this work, there has not been any evaluation yet, but we are planning to do a complete evaluation and here we outline how we will approach this. The evaluation would be done on two levels :

Macro level We will introduce the test subjects to the application, explain them the purpose, how it works and what are its functionalities. On this level we want to get answers to questions like : Is this application useful ? Does it address an actual need ? And if so, are the people aware of the existing need ?

Micro level In another evaluation, we focus more on the micro level. We want to know if the application is usable, which functionalities and features work well and which do not. In this evaluation the subjects would get specific tasks and we would then record how and how fast these tasks are completed. The specific tasks are not defined yet, but one example could be : Find the most interesting paper written by author x.

Public spaces In order to get more feedback, we also plan to deploy this visualization at one or more conferences, where we can observe the people discovering the tool and see what the initial thoughts are.

6 Future work

At the time of writing, a first working version of the application has been developed with some basic functionalities. But before we can do a real evaluation of this visualization, we need to improve the functionality of the application. In this section we describe the next steps that will be taken to achieve this.

An important feature that is missing at the moment, is being able to search for a certain author or paper to use as a starting point for the visual exploration. At a first stage we will add a keyboard like possibility to enter part of an author or a paper. To show the results there are several options that can be tried out. The found results can be highlighted in some way, or once a single result is found the visualization can center the result and zoom in on it.

At the moment it is not very visually clear yet which papers or authors are the most important or the most relevant. We are already exploring the possibilities to improve this by trying out filtering mechanisms and visual improvements. These visual improvements can be highlighting certain nodes or areas, varying the size of the nodes based on these factors, varying the strengths of the connections, etc.

7 Conclusion

In general, the fundamental issue is to understand in a deeper way how we can support the work of researchers with the technology that is available and how we can evaluate that our efforts make a difference. The design based research presented in this paper tries to move that agenda forward.

A major problem we face is getting clean data. At the moment this is too hard: we had to invest considerable effort in extracting the bibliographical data from the PDF version of the papers and in manually cleaning up the result. Initiatives like DBLP[11], Citeseer[12], bibsonomy[13], citeUlike[14] and others are targeting the same issue and we need to leverage their results in the context of our research.fm framework (see section 3.2) to create sustainable and scalable services for basic bibliographical data provision.

Assisting the user with navigation through the publication space is crucial. It is hard to figure out the correct way to combine navigation and search for manipulation of this information space. Currently, we only provide navigational access and we need to augment this with search facilities to locate relevant locations in this space: these can be papers or authors or relationships between them. We also need to add filtering facilities to reduce the complexity and size of this space to only that part that is relevant to the information need of an author.

We only use a fraction of the available metadata at the moment: our current visualization focuses on (co-)authorship relations between authors and papers. There is plenty of opportunity to also include other kinds of metadata in our scope: this could include forward and backward citations, geospatial information about the affiliations of the authors, textual relationships based on concept extraction techniques, etc. Assessing which kinds of such data help to address which kinds of problems researchers face and how we can exploit the data to make them useful and usable to that audience is a deep design challenge.

Finally, we do not exploit time information yet. However, especially as we start adding more of the metadata to our visualization, this will become an important concern. If we are able to integrate time information, then we can help users understand how a domain or publication outlet (conference, journal,

[11] http://www.informatik.uni-trier.de/ ley/db/
[12] http://citeseer.ist.psu.edu/
[13] http://www.bibsonomy.org/
[14] http://www.citeulike.org/

...) evolves, how a paper gains in influence, how the collaborative relationships between authors evolve, etc.

Acknowledgements We gratefully acknowledge the financial support of the European Commission through the STELLAR project and the STELLAR Network of Excellence.

References

1. P. Bergström and D.C. Atkinson. Augmenting the Exploration of Digital Libraries with Web-Based Visualizations. In *Fourth International Conference on Digital Information Management, 2009. ICDIM 2009*, pages 1–7. Citeseer, 2009.
2. V. Bush and A.W.M. Think. The Atlantic Monthly. *As We May Think*, 176(1):101–108, 1945.
3. R.A. Day. How to write and publish scientific papers. *Memórias do Instituto Oswaldo Cruz*, 93, 1998.
4. T. Dwyer, B. Lee, D. Fisher, K.I. Quinn, P. Isenberg, G. Robertson, and C. North. A Comparison of User-Generated and Automatic Graph Layouts. *IEEE Transactions on Visualization and Computer Graphics*, 15(6):961–968, 2009.
5. Clifton Forlines, Daniel Wigdor, Chia Shen, and Ravin Balakrishnan. Direct-touch vs. mouse input for tabletop displays. In *CHI '07: Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 647–656, New York, NY, USA, 2007. ACM.
6. Kenrick Kin, Maneesh Agrawala, and Tony DeRose. Determining the benefits of direct-touch, bimanual, and multifinger input on a multitouch workstation. In *GI '09: Proceedings of Graphics Interface 2009*, pages 119–124, Toronto, Ont., Canada, Canada, 2009. Canadian Information Processing Society.
7. J. Klerkx and E. Duval. Visualizing Social Bookmarks. In *conference Workshop on Social Information Retrieval for Technology-Enhanced Learning (SIRTEL07) at the EC-TEL conference. September*, pages 17–20. Citeseer, 2007.
8. J. Klerkx, E. Duval, and M. Meire. Using information visualization for accessing learning object repositories. In *International Conference on Information Visualisation*, volume 8, pages 465–470, 2004.
9. Xavier Ochoa, Gonzalo Mendez, and Erik Duval. Who we are: Analysis of 10 years of the ed-media conference. In *Proceedings of World Conference on Educational Multimedia, Hypermedia and Telecommunications ED-Media 2009*, pages 189–200, 2009.
10. Daniel M. Russell, Mark J. Stefik, Peter Pirolli, and Stuart K. Card. The cost structure of sensemaking. In *CHI '93: Proceedings of the INTERACT '93 and CHI '93 conference on Human factors in computing systems*, pages 269–276, New York, NY, USA, 1993. ACM.