

# SELECT: Social and Collaborative Filtering of Web Documents and News

The SELECT Project Team<sup>1</sup>

## Abstract

We describe the goals and progress to date of SELECT, a project funded under the EC Telematics Applications Programme. The objective of the SELECT project is to help Internet users to find the most reliable, valuable, important and interesting information quickly and easily and reduce information overload. SELECT is aimed not only at users who search for specific information, but also at those who use the Internet to keep up to date with what is happening in particular areas. In these ways, SELECT will make a positive contribution to the problem of helping users to tailor their information environments to meet their individual needs.

The approach adopted in SELECT is based upon information filtering. SELECT makes use of two filtering techniques. One is to make recommendations derived from an individual user's past choices. The other is to make recommendations derived from the behaviour of other users through social, collaborative filtering, especially those who have displayed similar tastes and interests in the past. Both techniques make use of users' ratings of Internet documents, either given explicitly or derived implicitly from evidence of users' behaviour.

## Introduction

The SELECT project is funded by the EC Telematics Applications Programme and involves partners from 9 countries. Its objective is to help Internet users to find the most reliable, valuable, important and interesting information quickly and easily, to avoid trash and reduce information overload. SELECT is aimed not only at users who search for specific information, but also at those who use the Internet to keep up to date with what is happening in particular areas. Project partners include one of the most successful European providers of Internet search services, and organisations representing more specialist user groups.

The Internet has opened up important new opportunities for knowledge exchange between scientific, technical, professional and other users. Sometimes, a user's need is to find particular information on particular topics, in other cases it is to update knowledge, to keep up-to-date with recent developments and increase contacts with other people with the same

---

<sup>1</sup>Roland Alton-Schiedl, TUV, Vienna, Austria.

Jesper Ekhall, Euroseek, Sweden.

Olivier van Geloven, Medibridge, Belgium.

Laszlo Kovacs and Andras Micsik, SZTAKI, Budapest, Hungary.

Christopher Lueg, IFI, Zurich, Switzerland.

Richard Messnarz, Independent Software Consulting Network, Ireland.

David Nichols, Computing Department, Lancaster University, UK.

Jacob Palme and Torgny Tholerus, KTH Technical University, Stockholm, Sweden.

Rob Procter and Dave Mason, Division of Informatics, University of Edinburgh, Scotland.

Enrico Stupazzini and Massimo Vassali, Omega Generation, Italy.

Richard Wheeler, StarLab, Brussels, Belgium.

Corresponding author: Rob Procter, Division of Informatics, University of Edinburgh, Kings Buildings, Edinburgh, EH9 3JZ, Scotland. Email: [rnp@dcs.ed.ac.uk](mailto:rnp@dcs.ed.ac.uk)

interests and specialities. There are two particular reasons why this is not easy to do. First, since anyone is allowed to put up any information they like on the Internet, there is no quality control (such as is done by editors and reviewers of magazines and journals). Second, there is often too much information, making it difficult to find what is most interesting and relevant.

It is clear that Internet users need tools so that they can find the information most valuable to them within the limited time they may have available. We believe that the free nature of the Internet is very important, however, so it is not our intention to implement techniques for censoring. Rather, our goal is to develop and implement techniques to aid users to find the information that is of the highest quality and relevance for their particular interests. In this way, SELECT will make a positive contribution to the problem of helping users to tailor their information environments to meet their individual needs.

SELECT is tackling this problem by using two main techniques. One is to make recommendations that are derived from an individual user's past choices. The other is to make recommendations derived from the behaviour of other users through social, collaborative filtering, especially those who have displayed similar tastes and interests in the past. Both techniques make use of users' ratings of Internet documents, either given explicitly or derived implicitly from evidence of users' behaviour.

The focus of SELECT is the WWW and Usenet News, the two most heavily used information domains on the Internet. The impact of the former is well known: it is estimated that several hundred million hours are collectively spent surfing the WWW per month. Usenet News is used by hundreds of thousands of people every day and generates an enormous amount of information. According to volume data published by the Swiss Academic and Research Network SWITCH, Usenet offers far more than 50,000 articles per day and the amount is increasing dramatically.

## **Information Filtering**

There has been considerable interest within the area of information filtering in recent times and several systems are now in use (see van Bommel, Koster and van der Weide (1997) for a good overview of ongoing research). For example, Sepia Technologies has developed a collaborative filtering system for movies, music and books.<sup>2</sup> Surflogic has developed Surfbot, a web browser plug-in that will search for and filter information on the net according to a user's needs.<sup>3</sup> The best known application of social filtering is Firefly, a commercial company that keeps a database of ratings of movies and music.<sup>4</sup> A user can connect, input his or her favourite movie or music, and be told which other movies and music were rated highly by people with similar tastes as the user. The MIT Centre for Coordination Science has developed GroupLens, a social filtering system for Usenet News (Resnick et al. 1994a, 1994b). Amazon, the Internet-based book retailer, offers its customers the opportunity to rate books and these are then offered to potential purchasers as a guide.<sup>5</sup>

We believe there are several reasons why the SELECT project will succeed better than previous projects. First, our intention is not to develop just another filtering system, but a filtering architecture into which different filtering methods can be plugged, compared and

---

<sup>2</sup> [http://www.sepia.com/suggestion\\_e.html](http://www.sepia.com/suggestion_e.html)

<sup>3</sup> <http://www.surflogic.com>

<sup>4</sup> <http://www.ffly.com>

<sup>5</sup> <http://www.amazon.com>

evaluated against one another. Most other projects in this area have just developed one single filtering method, and they have usually not attempted to perform any large-scale evaluation of user satisfaction with it. In SELECT, users will be able to test and evaluate different filtering methods against each other and indicate their opinions of them. Second, the SELECT project involves the creation of a ratings database for use in collaborative filtering. There have been some experiments before with collaborative filtering, but most have not been able to create the database needed. In SELECT, we have enlisted a major European Internet search service provider and their customers, which will enable the project to create a ratings database on a scale not previously achieved.

## **An Overview of SELECT**

The main focus of the SELECT project is the development of rating and filtering tools. By rating tools we mean mechanisms for users to evaluate and store their ratings of Internet documents and resources. By filtering tools we mean mechanisms to automatically scan Internet documents and resources before delivery to users. The result of filtering can be an ordering of documents and resources with the most interesting first, marking up of documents and resources with codes to help the user in manual decisions on what to read or even the discarding of less interesting documents.

Which of these approaches is used depends on users' needs. For areas important to a particular user, the filter will sort items, not discard them, but for less important areas, the user may prefer that the filter automatically discards documents. Filtering can also be based on ratings provided by the author of a document or by other readers of that document. In particular, the project is developing tools for social filtering, i.e., the selection of documents based on ratings made by people with similar values, competence and interests as the person for whom the filtering is done. This recommender group may be defined automatically by the system or users may nominate their own peer group. Users may belong to different recommender groups, reflecting their various distinct interests.

Part of the project is to develop and test different filtering methods: methods of assigning ratings to documents, methods of using these ratings to filter documents and methods of finding suitable filtering rules for each user. Automatic methods, where the system derives the filtering conditions from user actions or evaluations of documents and manual methods, where interaction between the user and the filter is used to establish the filtering conditions, are also being developed. At its simplest, however, SELECT simply allows users to perform the filtering themselves, informed by the ratings that others have given documents.

### **Rating**

Rating is implicit in most non Internet-based, traditional publishing services. For example:

- Newspapers, magazines, books, which are rated by their editors or publishers, selecting information that they think their readers will want.
- Consumer organisations and trade magazines which evaluate and rate products.
- Published reviews of books, music, theatre, films, etc.
- Peer review method of selecting submissions to scientific journals.

SELECT enables Internet users to input their evaluations of document quality. These are stored and used to aid the rater and/or other users to choose which documents to read. To improve the value of rater recommendations, the interest and knowledge profiles of users can be

matched to determine which users have with the same interests, values or knowledge. For example:

1. A user with a particular religion or political affiliation may prefer to find information which has been highly rated by other people with the same personal values.
2. A specialist in an area may want to find high quality information. Information that is of high quality for the specialist may be too complex for a beginner. Information that the specialist finds trivial may be valuable for a non-specialist who wants to learn the basics about a particular topic.

Using the tools provided by SELECT, rating may be applied to many kinds of documents, such as WWW pages, Usenet News postings, email, electronic journal papers, etc. Its purpose may be to increase the quality of the documents read, or to avoid certain documents deemed unsuitable in certain communities for certain groups of readers (for example: violence, pornography).

### **Filtering**

Filtering refers to the sifting of information according to some predefined criteria. This sifting has some similarities to information retrieval, but is also different in many aspects. In addition to the ratings given to a document by other readers (or by its author), filtering may be based on content-related criteria such as:

- Keywords in the document.
- Semantic analysis of the document.
- Analysis of the stylistic and genre qualities of the document.
- Analysis of the similarities between the document and other documents which the same user has rated highly.
- Documents directly related to other documents of high interest to the user, for example, by having hyperlinks to the document of interest.

Note that filtering is not only a matter of dividing all documents into two categories, good and bad, for a particular user. Often, what the user needs is instead a list of documents sorted by a matching index. Also, users may often want to sort information of interest into different areas representing their various interests.

Earlier studies of user requirements for rating and filtering tools have shown that different users have different requirements (Lantz 1993, Lantz 1995, Fahraeus 1997, Schmutzer et al. 1997, Irmay 1997). One user is interested in medieval religious beliefs; another is interested in particle physics. Another wants an overview of the knowledge on a certain topic; another wants to find the latest news. One user may want to get the maximum amount of information of value in a limited time, another user wants to browse and entertain at leisure. One user is an expert, another is a novice, in the subject area in which they are retrieving information.

The question is how can tools be designed to cater for all these differing users with their differing needs. However, even though users are different, they are common in that each user wants to find information of value to him or her. This is the basic user need that the SELECT project addresses. Its goal is not to find good information, according to some particular criterion of goodness, rather its aim is to develop tools that will make it easier for each Internet user to find the information that is important and interesting to them.

## Rating Sources

There are many different kinds of rating with different user requirements. In some domains, people are employed for making ratings. This is very common outside the Internet; most newspapers and journals have some rating system to decide what to publish and what to omit, even if they do not use this term for what they are doing. A special case is the peer review system used for choosing contributions to scholarly scientific and technical journals and conferences. In the electronic publishing area, this kind of rating is applied by portal services, perhaps the best known of which is Yahoo.<sup>6</sup> In Usenet News, moderated groups publish only contributions that have been approved by one or more moderators. A big disadvantage with human moderators is the delay they cause in publishing. In newsgroups and mailing lists, the time interval between one message and a reply to it is often only a few hours; in moderated lists, this time is lengthened to usually about a week. It is obvious that this can severely hinder rapid interaction in discussions.

		Right to rate a document	
		Everyone can input any rating (except some limitations on rating your own documents. This rule is probably not suitable to enforce automatically.	The right to input ratings is limited in some other way, to select people most proficient at providing good ratings in some way. Selection of such people may be a problem.
Use of ratings in filtering	Some kind of statistic (average, median, upper quartile) of all ratings set by everyone or by members of your peer group i.e., members of a professional organisation or expert in a particular area.	<p>Advantage: Lots of ratings available.</p> <p>Disadvantage: Ratings may not agree with your personal preferences.</p>	<p>Advantage: Better rating, may avoid misuse.</p> <p>Disadvantage: May reduce the amount of ratings available.</p>
	Ratings of people with similar views to yourself are preferably used through an automatic mechanism of comparing your ratings with those of other people.	Complex to implement but might provide very good ratings for your views and requirements. Encourages ratings, since only by giving ratings can your preferences be matched with others. Avoids problems of designating people with good competence to provide ratings.	Combines two different ways of trying to achieve the same thing: ratings set by those providing good ratings are given priority. This combination should not be used unless carefully analysed, since otherwise the two services can interact in unsuitable ways.

Table 1: Options for sources and uses of ratings.

Some systems follow an open rating principle, i.e., they allow anyone (or almost anyone) to rate any document. Sometimes just an average of these ratings is used, but some systems (e.g., Firefly) rate documents based on other people who have similar tastes (views, values, competence). A variant of this is to put people into different categories, so that users might specify that they prefer documents rated highly by other people in their own category (political or religious group, scientist, etc.). Document authors can also provide ratings, with

<sup>6</sup> <http://www.yahoo.com>

the advantage that more documents get rated, and that the ratings are easily transmitted with the document.

Table 1 summarises the interaction of two key dimensions of rating system design. The horizontal axis represent the options with respect to ratings submission, the vertical axis represents the options of whose ratings to use to guide the choice of documents.

After an initial investigation, it was decided that SELECT would focus on an open rating system. This raises some important issues. First, if anyone is allowed to submit ratings, there is a risk of misuse by people putting in high ratings on their own documents, or collusion between two people putting high ratings on their own documents. A check for the domain of the rater and the document can stop ratings by people in the same domain, but this is not a full protection. People known to misuse the rating system in this way can be identified and put on a stop list. Social codes that such misuse is not permitted may also help.

The second issue is how to get people to provide ratings. A good solution to this problem is that used by, e.g., Firefly, where users have to provide ratings to get access to the ratings of others. A variant of this is that a filtering system may use the ratings by a user as a tool in developing filtering conditions.

## **Explicit and Implicit Ratings**

Much of the current work on social filtering tools for the WWW has focused on so-called explicit methods, i.e., where the rater annotates a document (e.g., Bouthors and Dedie 1999) or (more simply) inputs a rating value. One drawback of this approach is that it calls for extra effort on the part of the rater, whilst failing to provide an equally immediate benefit, a recognised problem in collaborative systems (Grudin 1988). In contrast, implicit methods require no extra effort on the part of the rater, but have the disadvantage that the rating information provided has lower value. Some tools have attempted to find some middle ground between explicit and implicit approaches (Hill, Stead, Rosensteian and Furnas 1995). The SELECT project is pursuing the possibilities for using implicit, as well as explicit ratings, exploring how implicit approaches might be improved to provide rating information and higher value and relevance (Procter and McKinlay 1997).

In order to generate implicit ratings, it is necessary for users' behaviour to be observed. There are a number of kinds of information that can be extracted as side effects of users' browsing behaviour. These include:

1. Document read time. Morita and Shinoda (1994) reported a positive correlation between the time spent reading a document and the reader's assessment of its quality. The GroupLens project was able subsequently to verify this result for Usenet News postings (Konstan et al. 1997).
2. Documents that the user has bookmarked. However, surveys of Web users provides evidence that they typically bookmark fewer than 50% of the pages they find interesting; bookmarks tend to be evidence of strong, rather than marginal interest, so bookmarks set a relatively high threshold for recommendations (Rucker and Polanco 1997).

In addition, relevant information can be extracted from the documents themselves. For example:

1. Keywords, either as provided by the author or extracted automatically.
2. Text/image ratio, text/image hyperlink ratio and number of hyperlinks in the current document. From this information, it may be possible to derive some notion of genre and genre types into which documents may be categorised.
3. Language of the document.

In SELECT, users may also register a profile of their interests, likes and dislikes when they begin to use the service. This will be used in combination with the above data to generate an implicit rating of the document being viewed.

## **An Outline of SELECT System Architecture**

The SELECT architecture is based upon a client-server approach. The architecture is shown in Figure 1. The client side consists of a number of modules. The principal ones are as follows:

- User interfaces for ratings input and display.
- User interface for profile and preferences setting, e.g., interests, filtering rules, etc.
- Implicit rating module. This will work behind the scenes generating ratings derived from user behaviour and documents.

The principal server side modules are as follows:

- Passive filtering, i.e., where recommendations are simply based on the submitted ratings of documents.
- Active filtering, i.e., where recommendations are based upon the ratings of documents submitted by people with interests and/or rating histories that match those of the user. These recommender groups may be defined by the system from comparisons of previous behaviour, or may be user-nominated.
- A ratings database containing individual ratings submitted by both registered and anonymous users.
- A profiles database containing information about registered users, including their interests.

Client and server communicate using specially defined protocols. The functions supported include sending and requesting ratings, registering raters and exchanging ratings between SELECT servers.

## **An Example Implementation**

For a service like SELECT to be of greatest use, the availability of ratings derived recommendations needs to be as unrestricted as possible. That is, the system should be capable of making recommendations in whatever context a user happens upon a document. This means that the user is not required to explicitly ask for highly recommended documents (say on a particular topic), but that the user will be able to see the SELECT service's recommendations seamlessly integrated with the circumstances in which the documents become available.

As an illustration, consider the WWW page that a user is currently reading. This page may contain links to other pages or documents. The user's problem is to choose which of these

links to follow. In other words, the user needs SELECT's recommendations pertaining to these (as yet unseen) documents. One solution which satisfies the requirement for seamless access



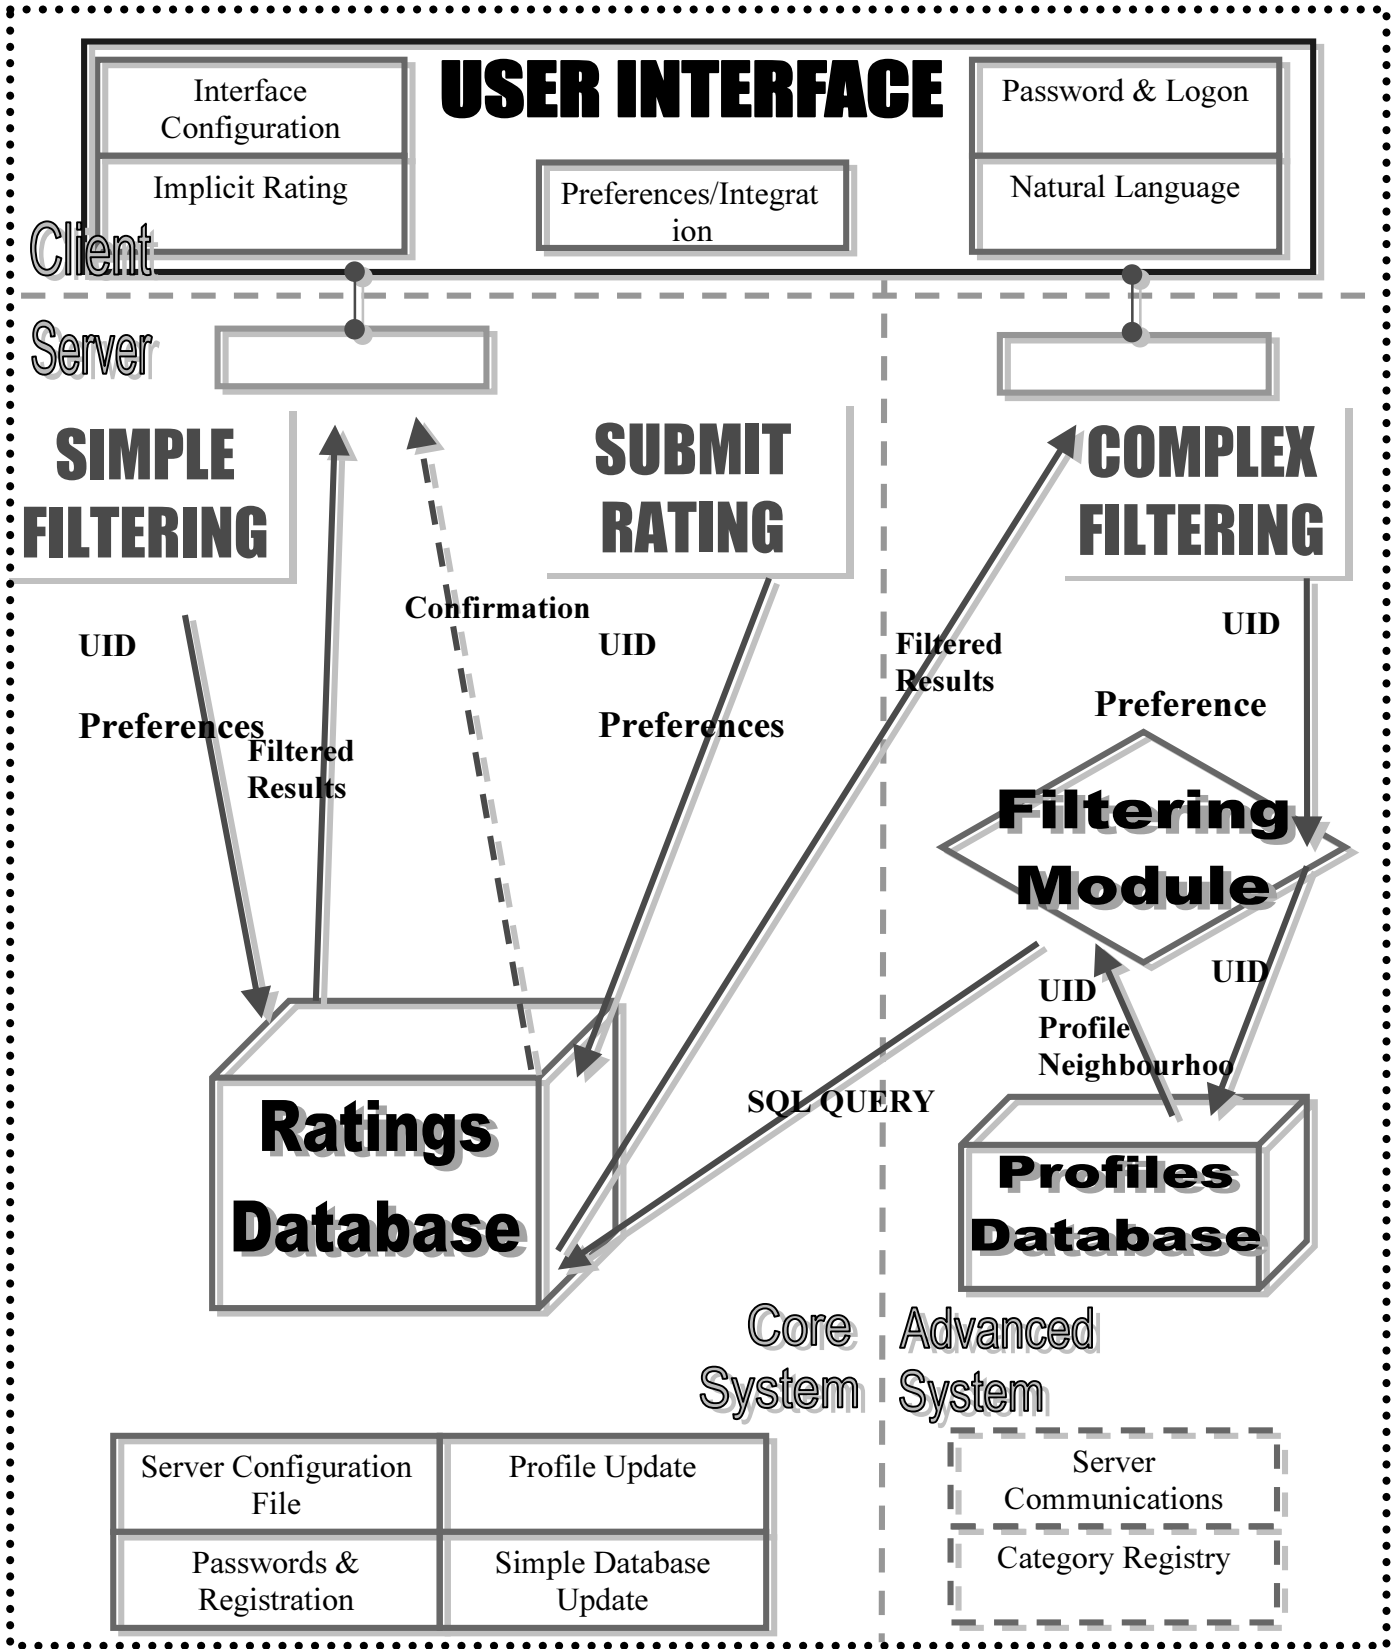


Figure 1: SELECT system architecture.

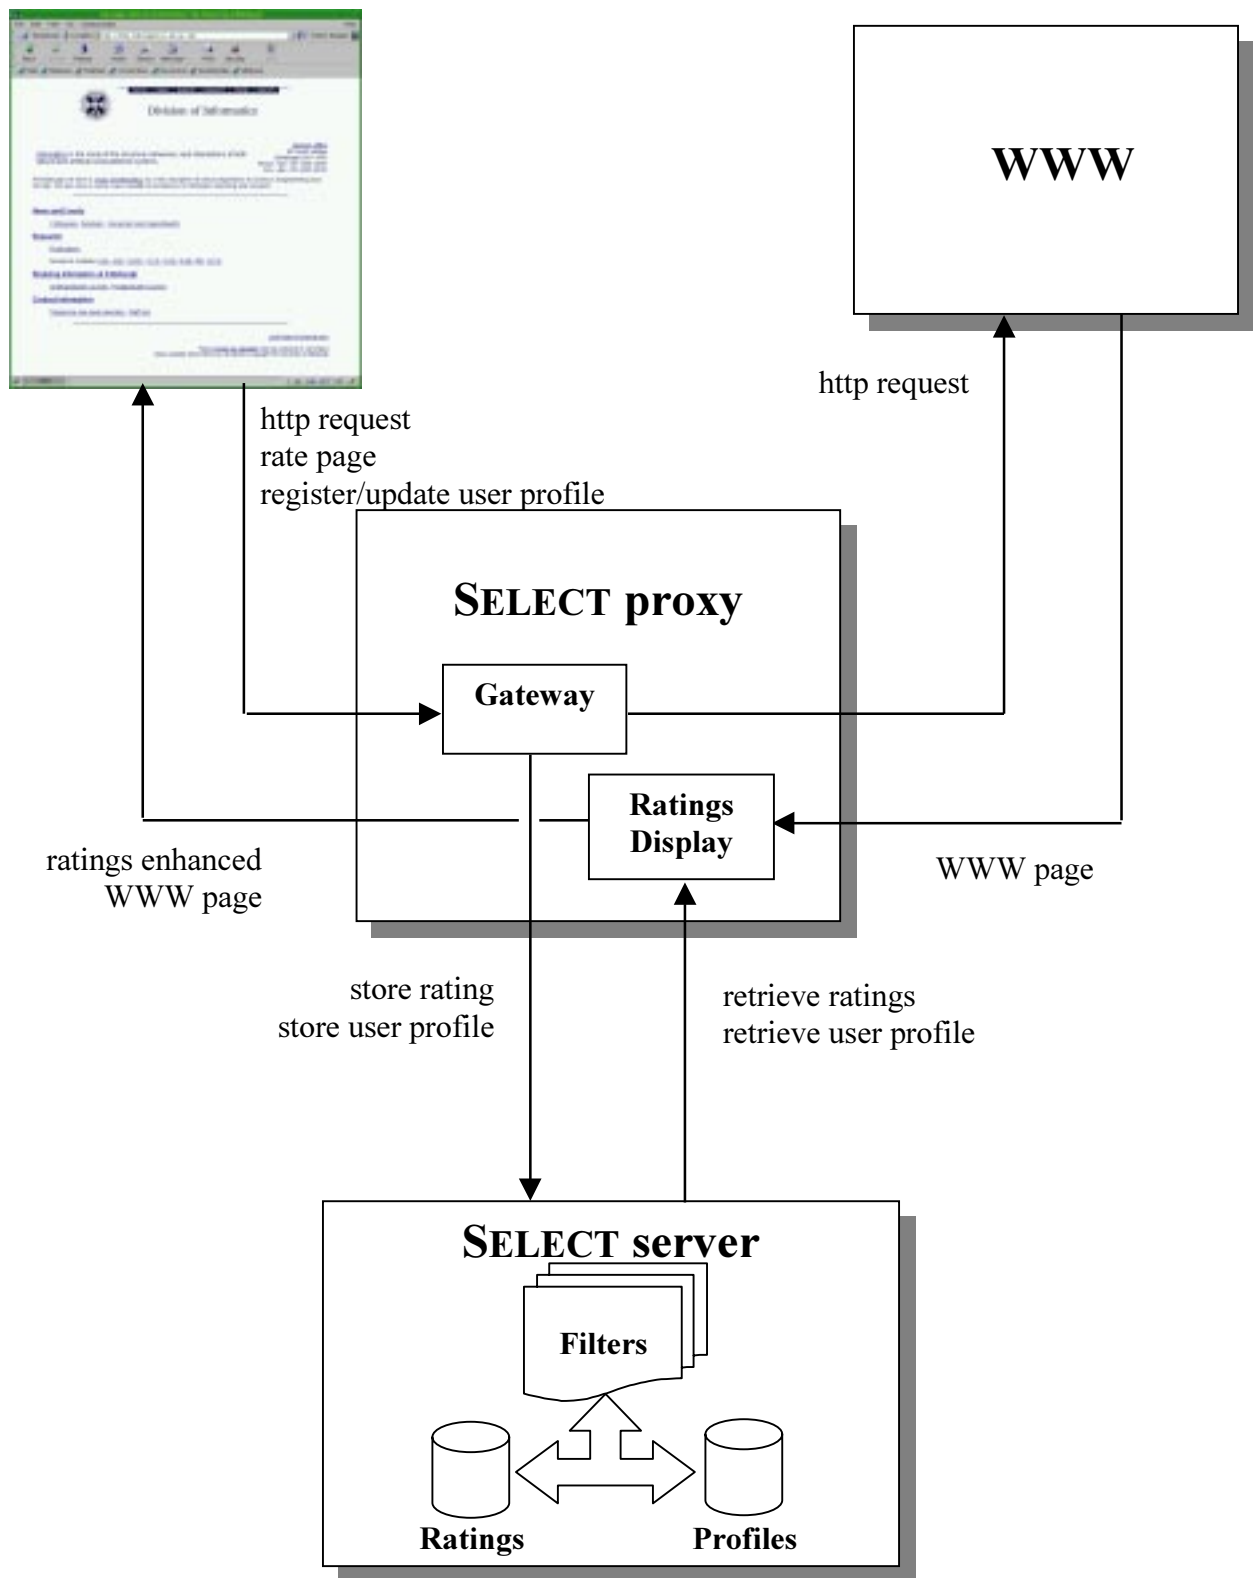


Figure 2: A proxy-based implementation of SELECT.

is to use a proxy-based approach, similar to that used, for example, in WebWatcher (Joachims, Frietag and Mitchell 1996) and Pharos (Bouthors and Dedieu 1999).

In general, a proxy-based approach allows the content of a document, say a WWW page, to be modified in a variety of ways. Content may be modified or removed before delivery to the user; there are several proxies that allow advertisements to be removed from WWW pages (e.g., Intermute<sup>7</sup>). Alternatively, documents can be enhanced with ratings or annotations (e.g., CritSuite<sup>8</sup>). In addition to single-purpose client proxies there are several existing proxies which can be extended to perform arbitrary modification of WWW content (e.g., Muffin<sup>9</sup>). In combination with a ratings server such a proxy can enhance WWW pages to present the user with social feedback on the quality of links in the current document.

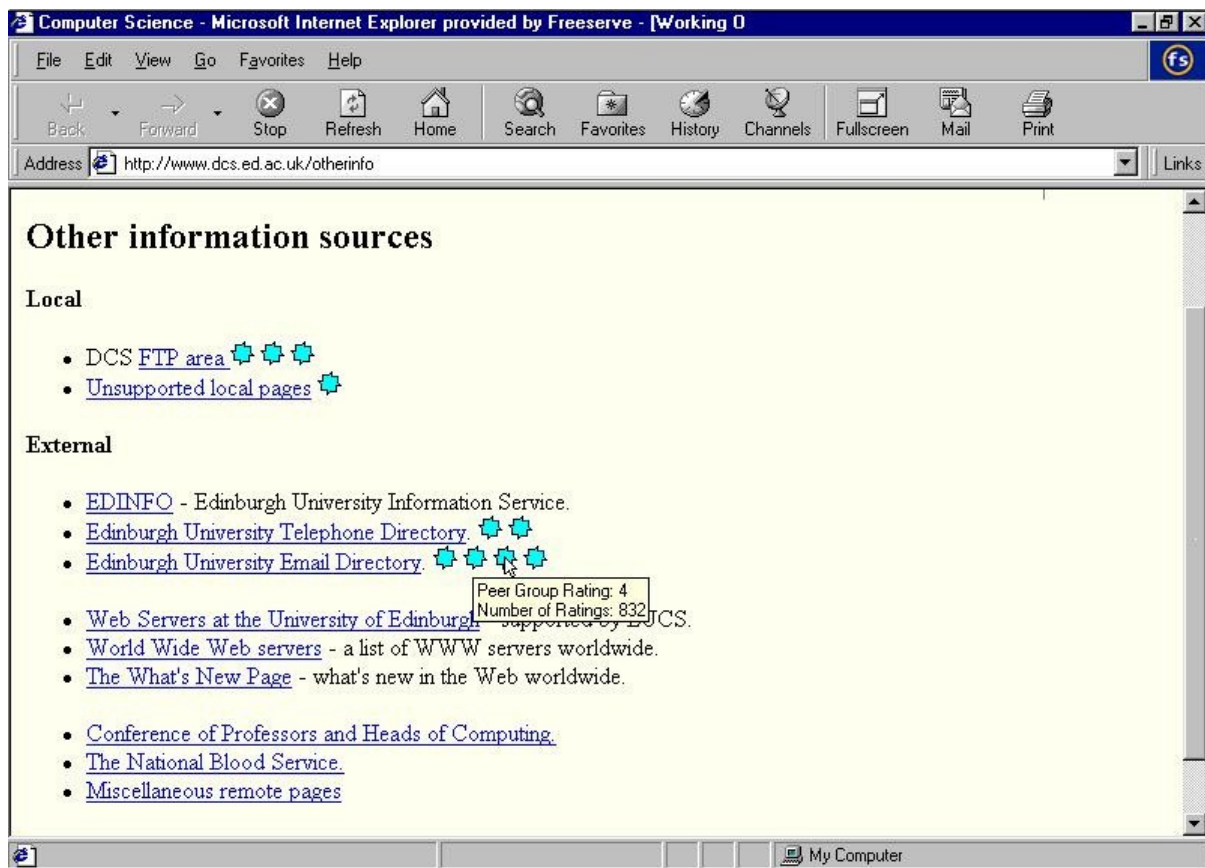


Figure 3: An example ratings-enhanced WWW page.

A proxy-based approach is sketched in Figure 2. As each WWW page request is made by the user, the proxy sends a query to the SELECT server requesting the ratings that are held for this particular document. When the WWW page is returned by its host site, relevant ratings information are returned by the proxy and are displayed on the user's desktop. More interesting, however, are the options that a proxy-based implementation opens up for enhancing information about the links contained within the WWW page. An example is shown in Figure 3. Here, ratings for individual links are indicated by the '\*' symbols. Passing the cursor over the rating causes a floating pop-up window to appear with more details about

<sup>7</sup> <http://www.intermute.com/>

<sup>8</sup> Foresight Institute, <http://crit.org/>

<sup>9</sup> <http://muffin.doit.org/>

how the rating was been derived. In the example, this indicates that it has been calculated from over 800 individual ratings submitted by members of that user's peer group.

Other services, such as site statistics, download predictions and document previewing may also be incorporated within the same basic architecture (Stanyer and Procter 1999).

## User Interfaces for Ratings Input and Display

It is critical that the effort required to make an explicit rating be minimised. At the same time, it is also important that raters provide enough information for the rating to be capable of being interpreted accurately. The GroupLens project reduced user inputs to a single click by combining the rating function with the 'next article' function. Although this does not seem feasible in a WWW domain (where the user may follow any link on a page), it is indicative of the sensitivity of collaborative filtering systems to rating input costs. A lack of ratings may prove especially significant at the start of a collaborative system: where the functionality is dependent on ratings input by users then the system has few benefits to a potential user (Lueg 1998). This 'cold-start' problem is usually addressed by using content-based techniques, e.g., filterbots (Sarwar et al. 1998). In SELECT we can also consider using implicit techniques for a period before introducing explicit ratings.




Figure 4: A prototype ratings input user interface.

Even when ratings are being provided and used to recommend resources, there may still be dis-incentives for users to provide ratings (Avery and Zeckhauser 1997). Although the user interface may enforce a 'stick' approach (not providing rated documents to those who do not contribute ratings) it can also enhance incentives for rating by using social 'carrots'. Under

certain circumstances users will perform ostensibly altruistically, i.e. they may provide ratings without tangible benefits. One such intangible benefit could be 'fame' -- the user interface could show who is contributing to the community by providing ratings, or show that a particular recommendation has been derived mainly from a small group of raters.

Although it is in users' interests to provide ratings, as they help personalise any recommendations, by providing an alternative 'benefit' the SELECT user interface may be able to achieve the necessary 'critical mass' of users. The visibility of such social benefits helps balance the (possibly) 'black box' benefits of increased personalisation. In order to be part of a successful collaborative filtering system the interface then has to not only be sensitive to the cognitive costs of ratings input, but also support wider social relations. The user interface is not merely a mechanism for connecting users to the ratings database but it reflects the organisation of the virtual community and the contributions (or otherwise) of the participants. The Knowledge Pump (Glance, Arregui and Dardenne 1998) is an example of a user interface that attempts to reflect these concerns within an intranet environment; a challenge for SELECT is to generalise to the environment of the Internet.

A survey was carried out to determine which of several different ratings input user interfaces people found most acceptable. The most popular choice is shown in Figure 4. This example uses a frame-based approach, the frame being added by the proxy as above. The rating input interface incorporates a display of the WWW page's current rating, shown here as both an average and a population distribution for the individual ratings.

## Summary

The goal of the SELECT project is to develop and implement techniques to aid users to find the information that is of the highest quality and relevance for their particular interests. In this way, SELECT will make a positive contribution to the problem of helping users to tailor their information environments to meet their individual needs.

The approach adopted by is SELECT based upon two main techniques. The first is to make recommendations that are derived from an individual user's past choices. The second is to make recommendations derived from the behaviour of other users through social, collaborative filtering, especially those who have displayed similar tastes and interests in the past. Both approaches make use of users' ratings of Internet documents, either given explicitly or derived implicitly from evidence of users' behaviour. More than simply developing specific implementations, however, the project aims to create an information filtering architecture that will afford the use of new techniques as and when they emerge.

Implementation of the base system is well under way and it is expected to be ready for testing and evaluation by user groups by the end of the year.

## References

- Avery, C. and Zeckhauser, R. Recommender systems for evaluating computer messages. *Communications of the ACM*, vol.40, no.3, March, 1997, p. 88-9.
- Bouthors, V. and Dedie, O. Pharos, a Collaborative Infrastructure for Web Knowledge Sharing. Research Report RR-3679, 1999. INRIA.
- Fåhræus, E. Intelligent Filtering on Usenet News: A User Study. Technical Report 97-003, Department of Computer and Systems Sciences, Stockholm University/KTH, 1997. <http://www.dsv.su.se/~evafaahr/Filter.html>

- Glance, N., Arregui, D. and Dardenne, M. Knowledge Pump: community centered collaborative filtering. In Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering, Budapest, Hungary, 1998, p. 83-88. ERCIM Press.
- Grudin, J. Why CSCW applications fail: problems in the design and evaluation of organisational interfaces. In Proceedings of CSCW'88, New York, 1988, p. 85-93. ACM Press.
- Hill, W., Stead, L. Rosensteian, M. and Furnas, G. Recommending and Evaluating Choices in a Virtual Community of Use. In Proceedings of CHI'95, Denver, 1995. ACM Press.
- Irmay, M. Some reflections about information filtering on the Internet, Swiss Federal Research Institute, 1997.
- Joachims, T., Freitag, D. and Mitchell, T. WebWatcher: A Tour Guide for the World Wide Web. Research Report CMU-CS-xxx, School of Computer Science, Carnegie Mellon University, Pittsburgh, USA, 1996.
- Konstan, J., Miller, B., Maltz, D., Herlocker, J., Gordon, L. and Riedl, J. GroupLens: Collaborative Filtering for Usenet News. Communications of the ACM, March, 1997, p. 77-87.
- Lantz, A. How do experienced Users of the System Usenet News select their Information? Technical report, Department of Computer and Systems Sciences, Stockholm University/KTH, 1993.
- Lantz, A. Useful Criteria for Intelligent Filtering? Technical report 95-042, Department of Computer and Systems Sciences, Stockholm University/KTH, 1995.
- Lueg, C. Considering Collaborative Filtering as Groupware: Experiences and Lessons Learned. In Reimer, U. (Ed.) Proceedings of the Second International Conference on Practical Applications of Knowledge Management (PAKM'98), October, 1998.
- Morita, M. and Shinoda, Y. Information Filtering Based on User Behavior Analysis and Best Match Text Retrieval. In Proceedings of the ACM-SIGIR Conference on Research and Development in Information Retrieval (SIGIR'94), 1994.
- Procter, R. and McKinlay, A. Social Affordances and Implicit Ratings for Social Filtering on the Web. In Proceedings of the Fifth DELOS Workshop on Filtering and Collaborative Filtering, Budapest, 1997. ERCIM Press, p. 89-96.
- Resnick, P. GroupLens: An Open Architecture for Collaborative Filtering of Netnews. In Proceedings of CSCW'94, Chapel Hill, 1994, ACM Press, p. 175-186.
- Resnick, P., Zeckhauser, R. and Avery, C. Roles for Electronic Brokers, Twenty-Second Annual Telecommunications Policy Research Conference, October, 1994.
- Rucker, J. and Polanco, M. Personalized Navigation for the Web. Communications of the ACM, March, 1997. p. 73-75.
- Sarwar, B.M., Konstan, J., Borchers, A., Herlocker, J., Miller, B. and Riedl, J. Using filtering agents to improve prediction quality in the GroupLens Research collaborative filtering system. In Proceedings CSCW'98, New York, 1998. p. 345-54. ACM Press.
- Schmutzer, R. Scenarios for Voting and Rating Using Web4Groups. In R. Alton-Schiedl, R. Schmutzer, P. Sint and G. Tscherteu (Eds.) Rating, Voting and Annotations. Oldenberg, 1997.
- Stanyer, D. and Procter, R. Improving Web Usability with the Link Lens. In Mendelzon, A. et al. (Eds.), Journal of Computer Networks and ISDN Systems, Vol. 31, Proceedings of the Eighth International WWW Conference, Toronto, 1999. Elsevier, p. 455-66.
- van Bommel, P., Koster, C. and van der Weide, Th. Profile – A Proactive Information Filter. Technical Note CSI-N9602, Computing Science Institute, University of Nijmegen, 1996.

Cite as:

Alton-Scheidl, R., Ekhal, J., van Geloven, O. Kovacs, L., Micsik, A., Lueg, C., Messnarz, R., Nichols, D., Palme, J., Tholerus, T., Mason, D., Proctor, R., Stupazzini, E., Vassali, M., and Wheeler, R. (1999) SELECT: Social and collaborative filtering of web documents and news, *Proceedings of the 5th ERCIM Workshop on User Interfaces for All: User-Tailored Information Environments*, Dagstuhl, Germany. ERCIM. 23-37.