

Who controlled distribution of information before the Internet?

- Publishers
- Newspapers
- Organisations (societies, companies, political parties, etc.)
- Schools and universities
- Governments

Was there no free speech then?

Sure, there was free speech. But free speech was controlled, channelled and organised.

In whose interest was it controlled, channelled and organised?

- Politicians
- Establishment
- Scientific community
- Readers

What is different with the Internet?

- Anyone can easily at low cost publish anything they want.

Is this good or bad?

Both!

But everyone cannot read everything. The systems for control of free speech which we had before the Internet were in many ways tools which aided people in selecting the most valuable information.

Newspapers and magazines selected the most interesting news. You chose to subscribe to the paper which selected according to your tastes.

The same with television, books.

The same with societies: You chose which society to join, and in that way selected what information you would get and could yourself disseminate.

The quality of the information on the Internet is very varying. There is lots of interesting things, but also lots of trash. (Not that everyone agrees on what is interesting and what is trash, of course.)

What is filtering?

Filtering is tools to help you find the most valuable information, so that the limited time you have for reading can be spent reading interesting information and avoiding trash.

Automatic filtering is where the computer evaluates what is of value for you.

Social filtering is tools where other people help you evaluate what is of most value to read. Just like the publishers and organisations did in society before the Internet.

The most successful social filtering system is Yahoo. Yahoo employs humans to evaluate documents, and puts documents which are interesting into its structured information data base. Just like the publishers and organisations did in the world before the Internet.

Is filtering successful?

Automatic filtering is successful only with very simple filters.

Examples of partially successful filters:

- Filtering by mailing list/newsgroup.
- Filtering by topic (thread).

But filters which automatically in a more intelligent way finds what is of most interest to you have not been very successful.

Why?

Maybe because filtering is a complex task requiring intelligence which computers are not yet capable of?

The most successful social filtering service, Yahoo, employs people to select the best, just like publishers did before the Internet.

Another important social filtering service is mailing lists, where only members can contribute (or at least should be able to) and various tools are used to keep out disruptive people.

How can social filtering work on the Internet?

Several different ways:

1. People are employed to make the selection (like Yahoo).
2. More or less closed groups for people with common interests.
3. Peer groups helping each other find the best.
4. Data bases where anyone can put their ratings:
 - a) You get the highest rated documents.
 - b) You get the highest rated documents, as rated by those who make similar ratings as you do.

In case 4b above, how do you find people with similar interests, views and values as yourself?

- A. By choosing predefined categories (joining groups, societies).
- B. By automatic derivation of categories based on the ratings you supply yourself.

Compare 1-3 and A, More manual, with 4 and B, more automatic.
What do people want?

How is research on filtering usually done?

A clever computer scientist develops his wonderful new ideas of how filters should work. His/her results very seldom result in products used by other people.

Is there another way of organising research?

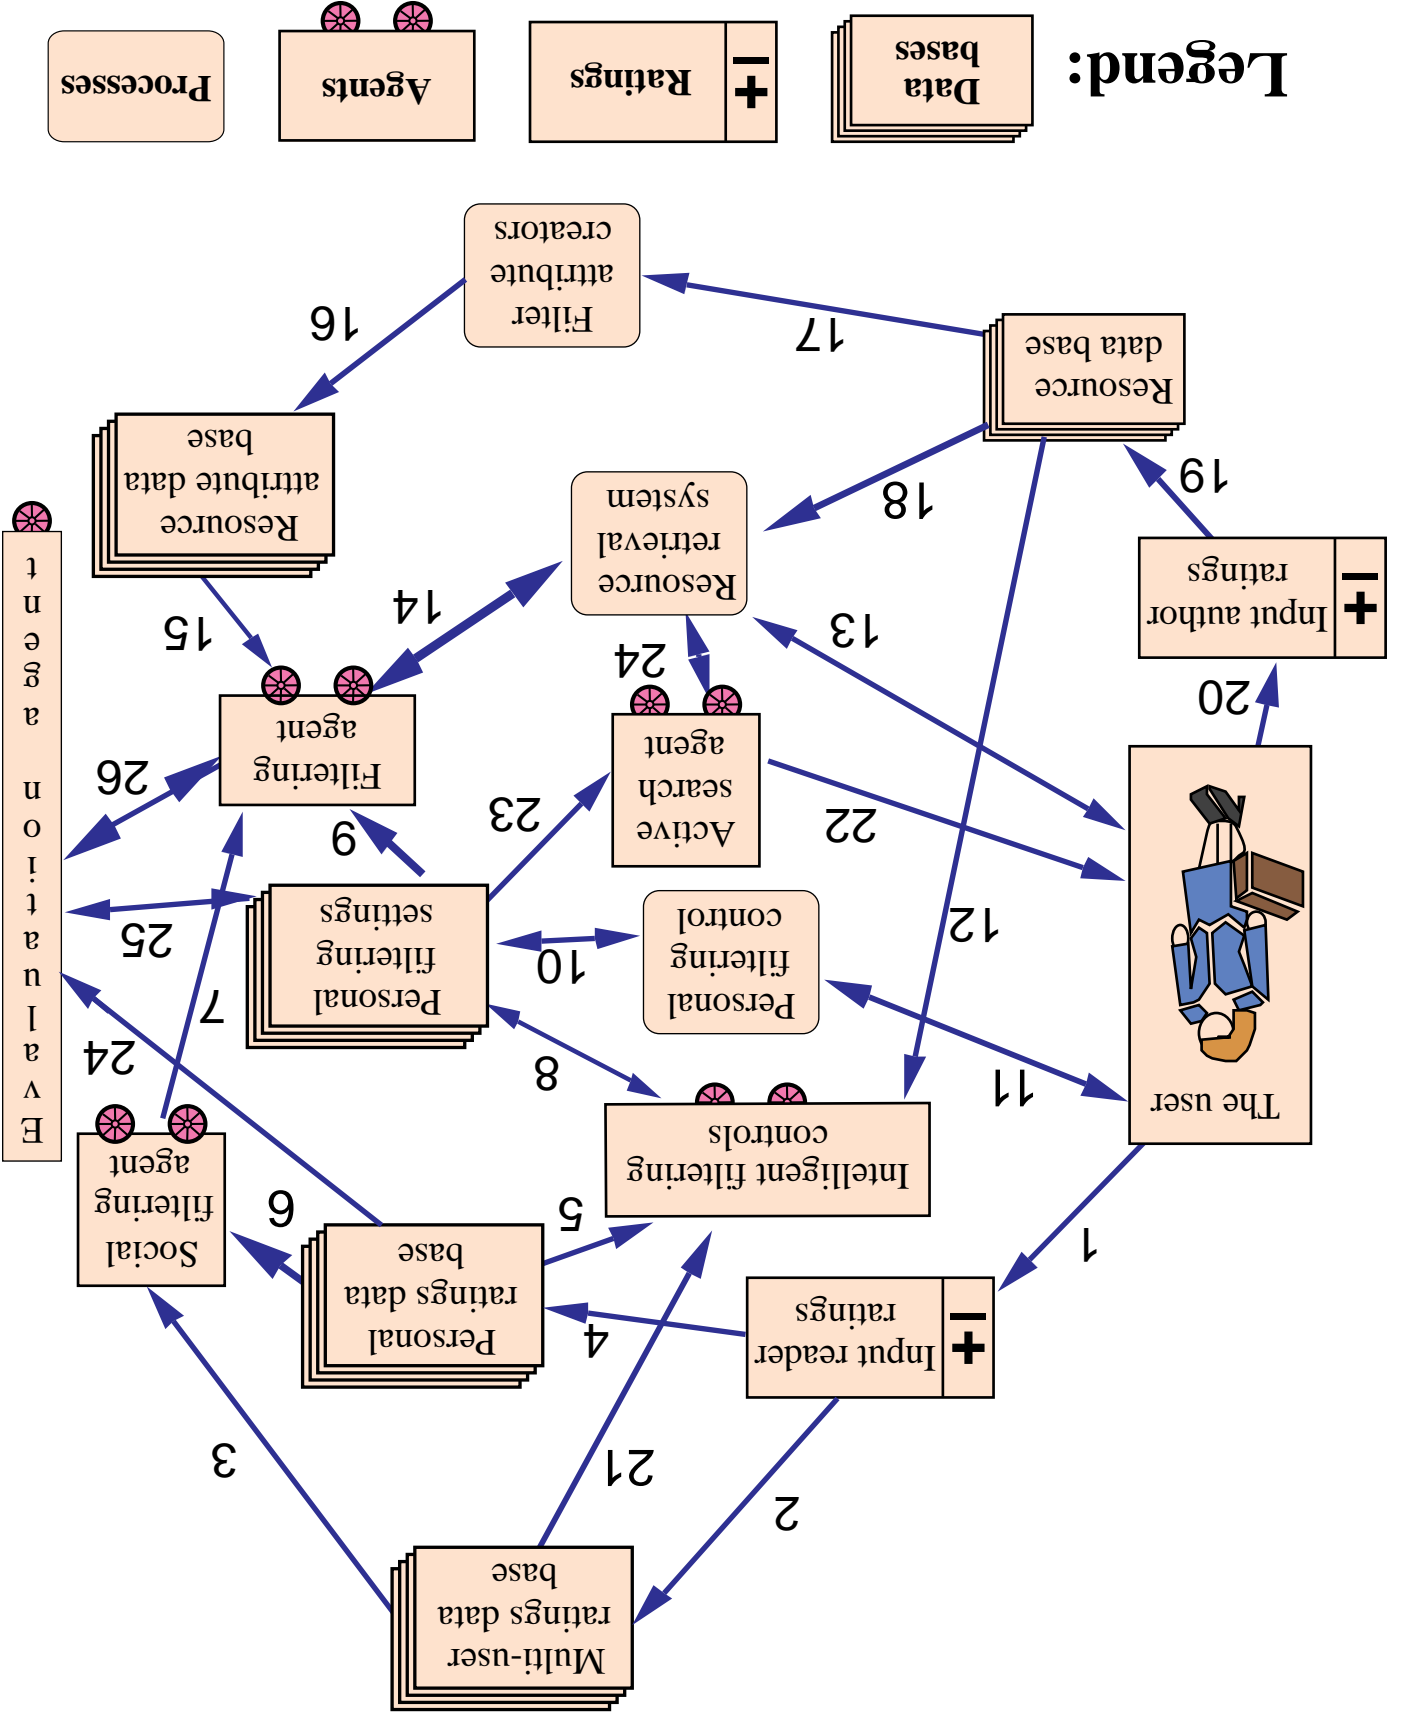
Yes: Specify a general-purpose architecture, so that many different developers can develop and test their filters.

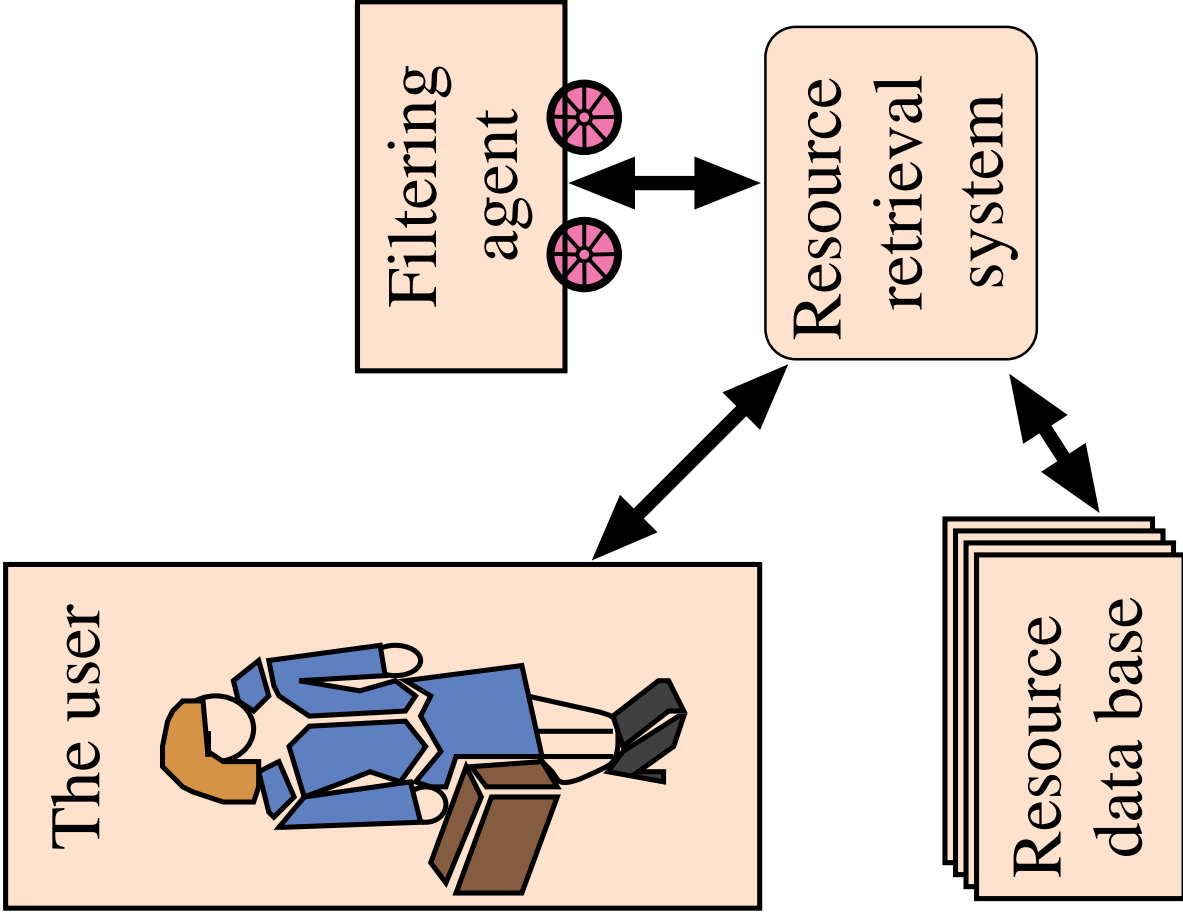
An architecture is a defined set of modules and interfaces, where there can be different versions of each module, which can still co-work with the other modules.

Why should we specify such an architecture?

1. Developers need not develop the whole filtering system, with all its components, which usually is very time-consuming. They can develop only the module where they have bright ideas, and these modules can be tested in the general-purpose architecture.
2. Different filtering methods and modules can be tested and evaluated against each other within the general-purpose architecture.

Proposed architecture for multi-user filtering and rating system for SELECT project, November 1997 by Jacob Palme.






Resource retrieval systems and data bases

- **The Web**
- **Usenet news**
- **E-mail with mailing lists**
- **Other, like Web4Groups**

Major Components

- **Resource storage (Distribution, replication, archiving)**
- **Resource structuring (Search engines, Subject trees)**
- **News control**
- **User interface**



Filtering agent, settings, rules, resources

Applies personal filtering settings and rules to resources and resource attributes.

Filtering rules: Usually collections of boolean expressions.

Resource attribute data base: Derived attributes of resources, can be keywords, genre, quality, etc.

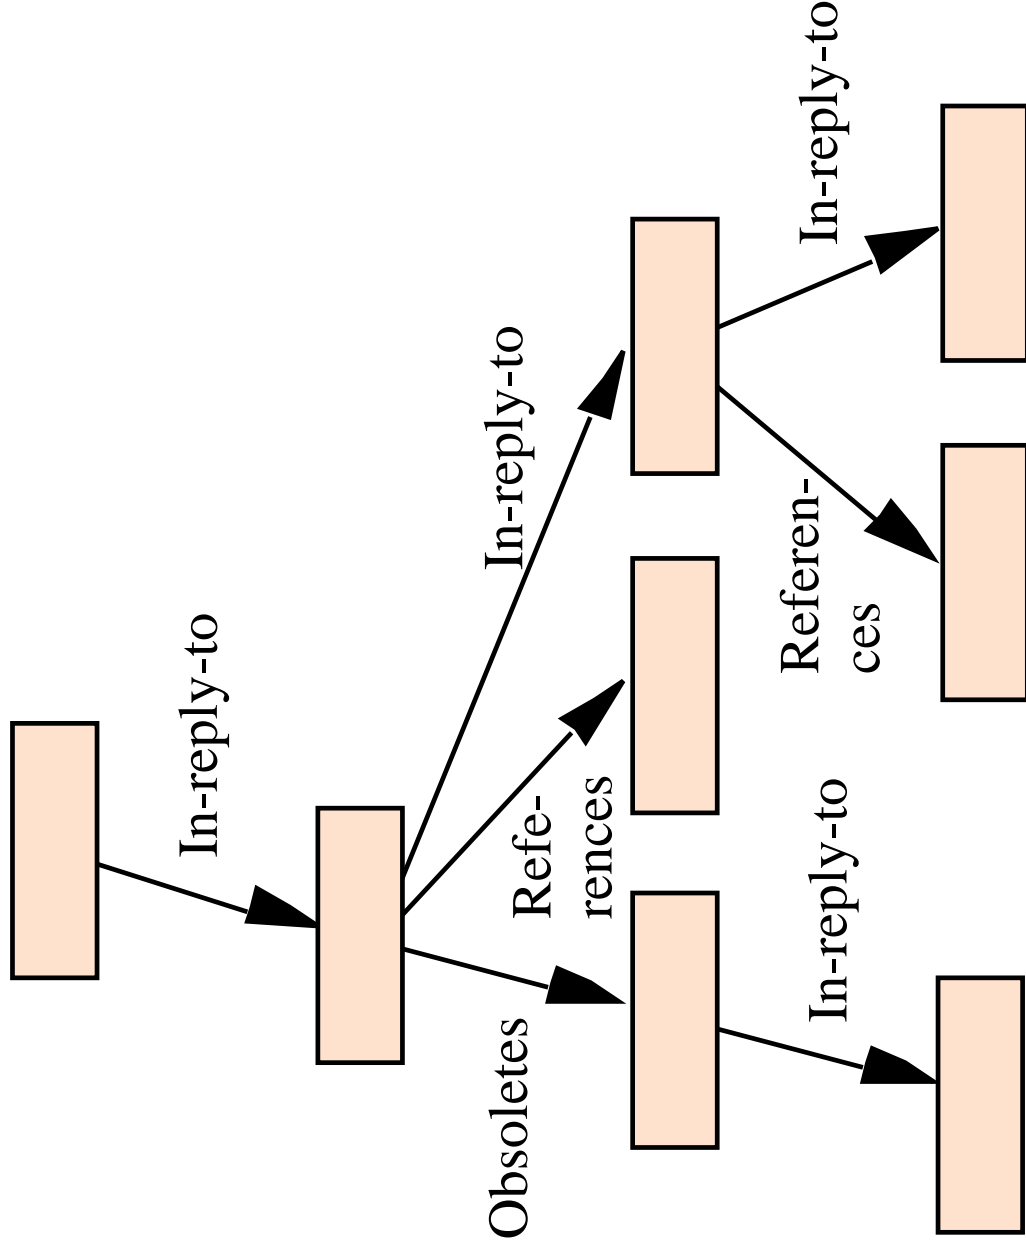
Issues

Should filtering agent be placed in server or client? Operate at download or in the background?

Syntax and semantics of filtering rules. Sequential execution?

Derive keywords in advance? Which resource attributes?

Example of a thread in e-mail




HTML hyperlinks can also form threads.

Filtering of threads
Filtering should not be done on small single documents. Such documents may get new importance when linked to a set of documents, such as discussion threads or hyperlink structures.

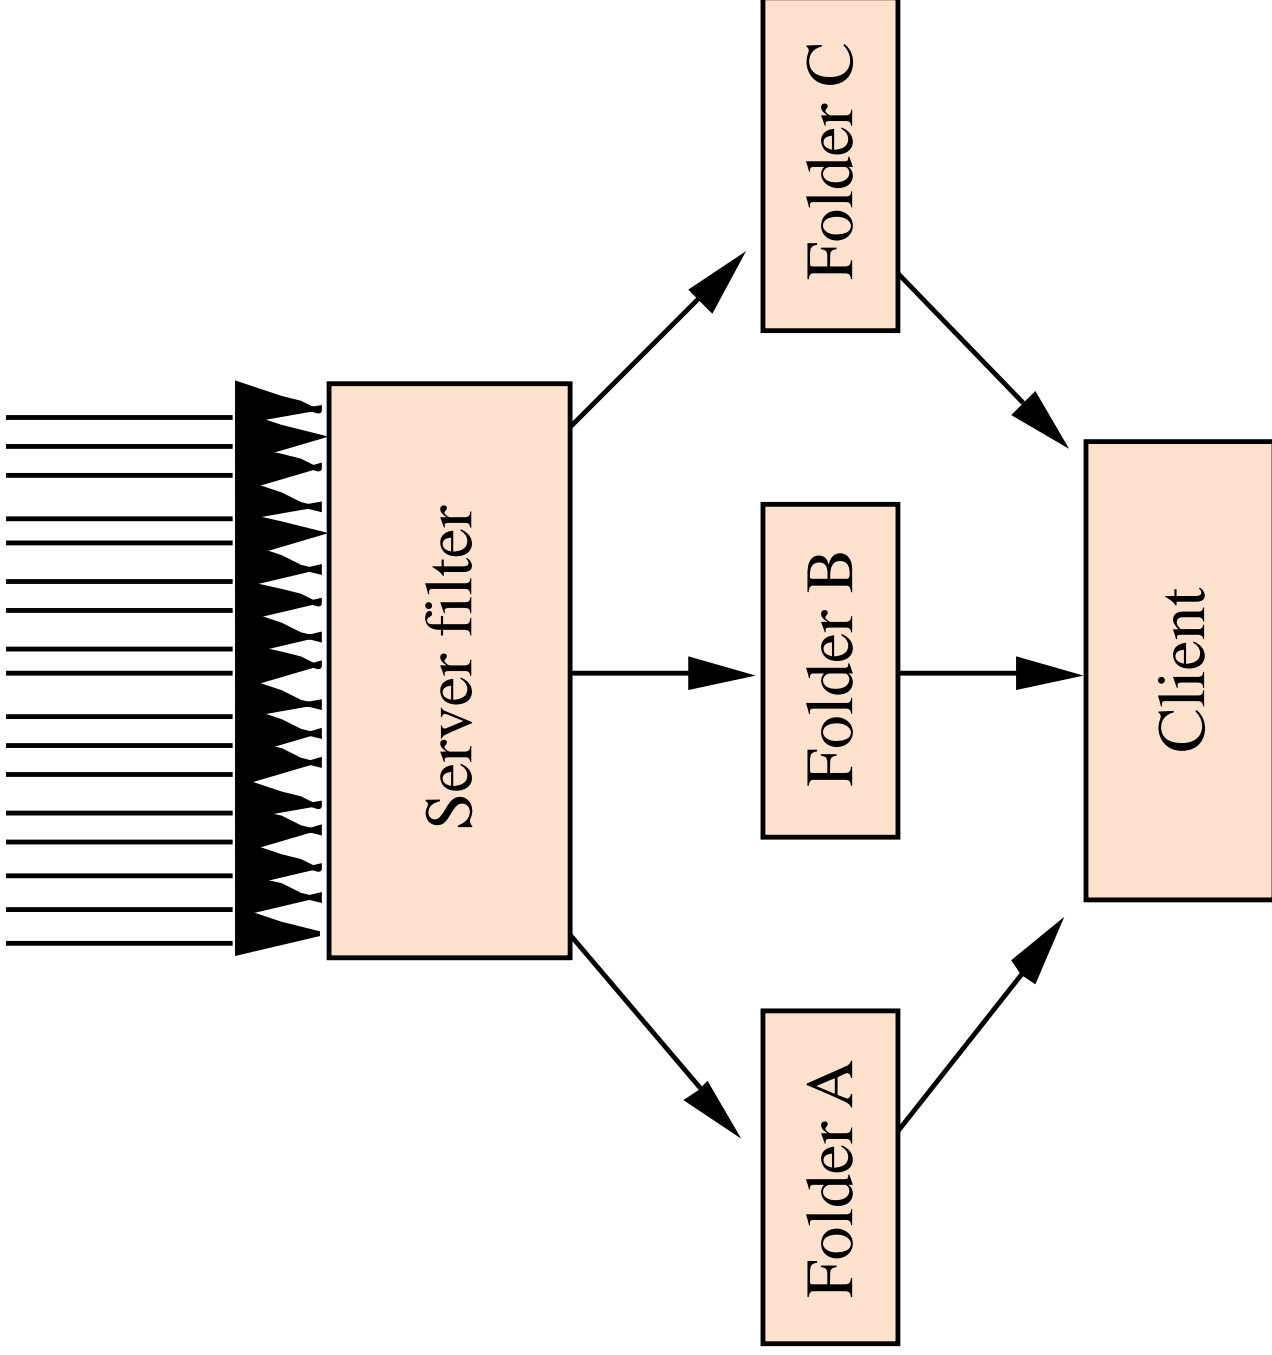
Difficult to achieve if filtering agent examines each document separately.

Common feature in Usenet News; so-called kill buffers.



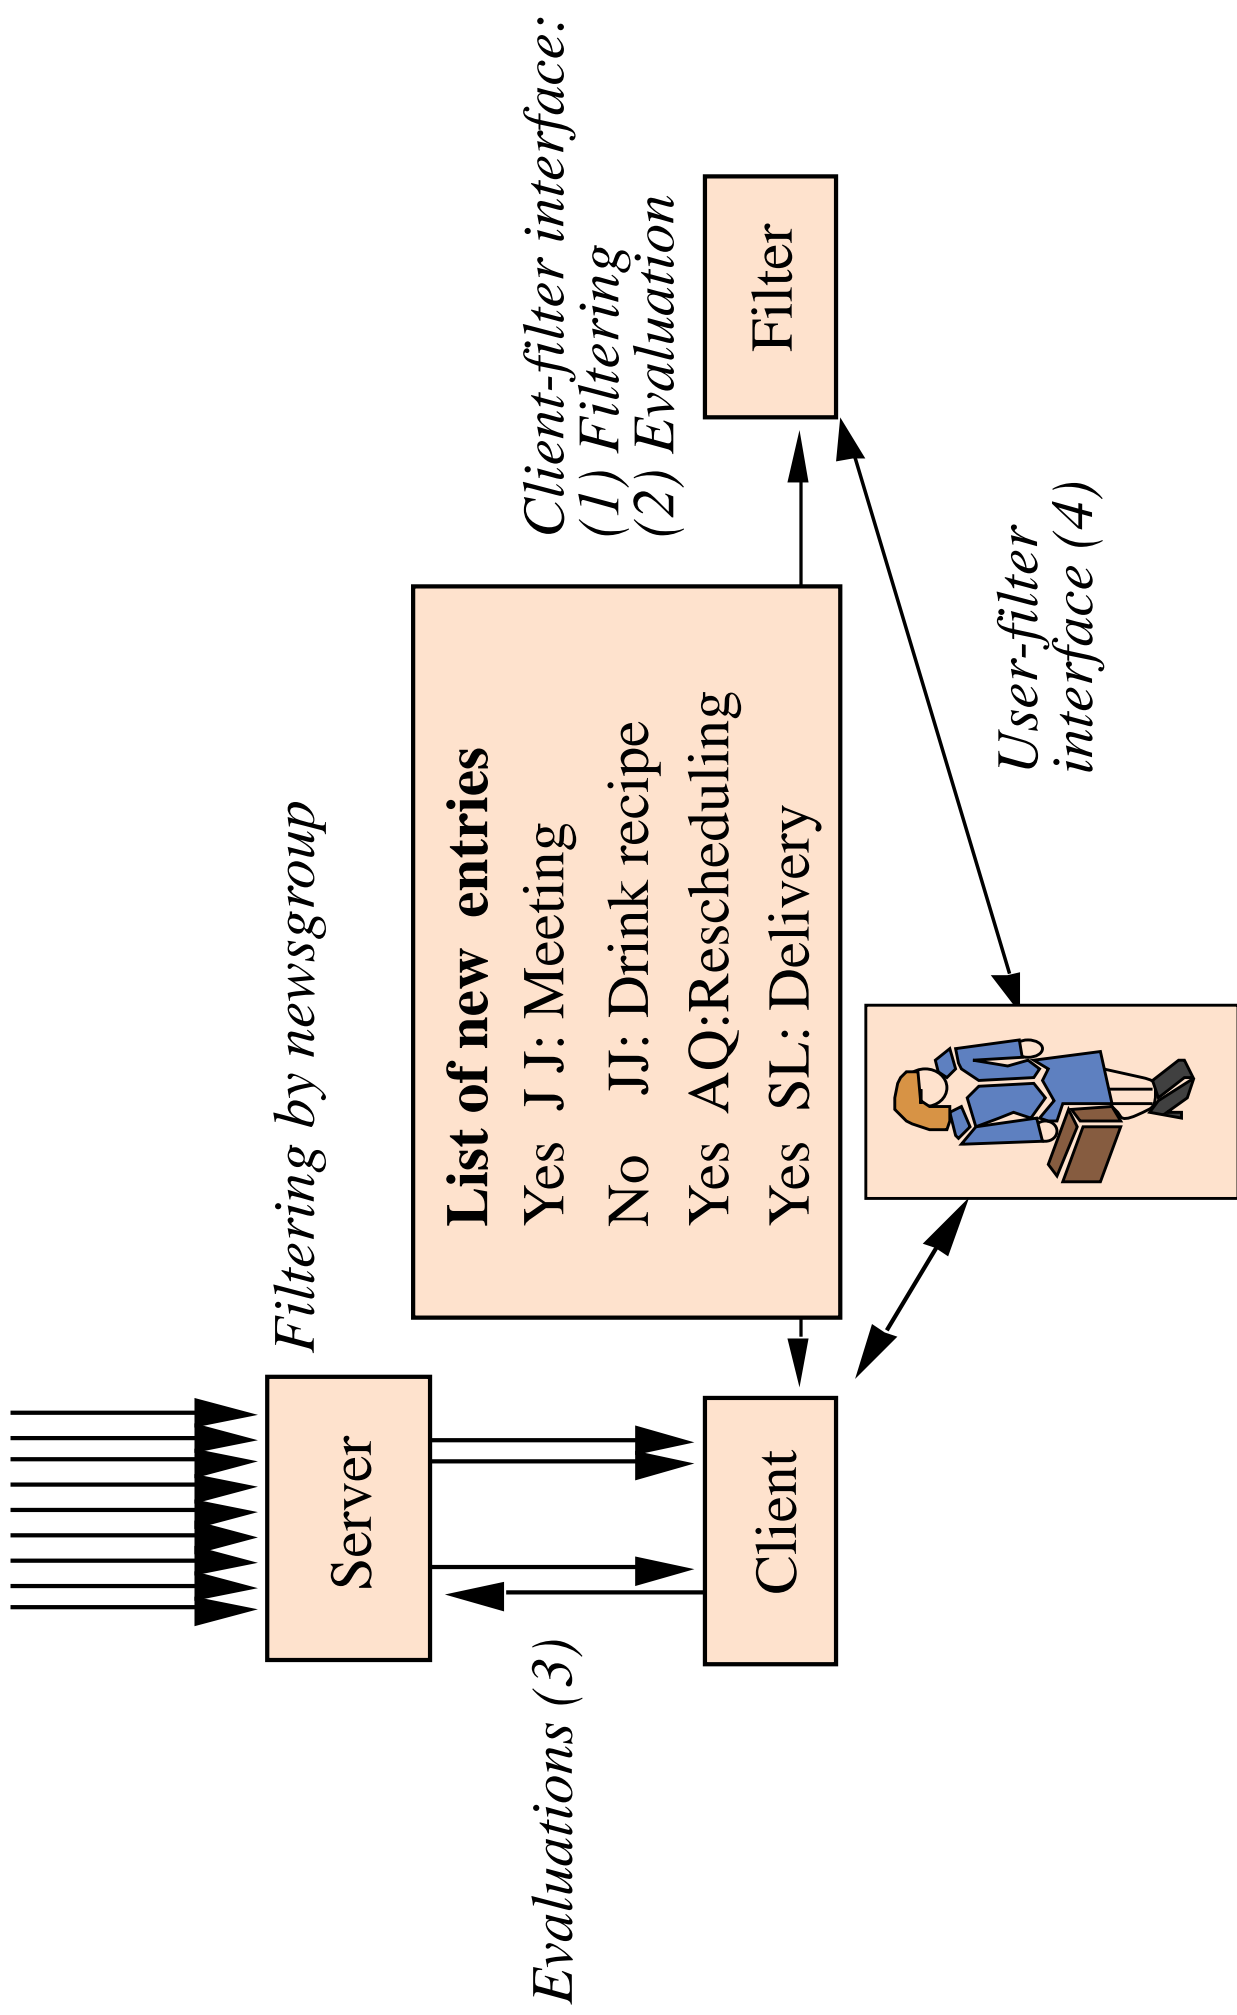
Typical messaging architecture

Where should we put a filter in this architecture?

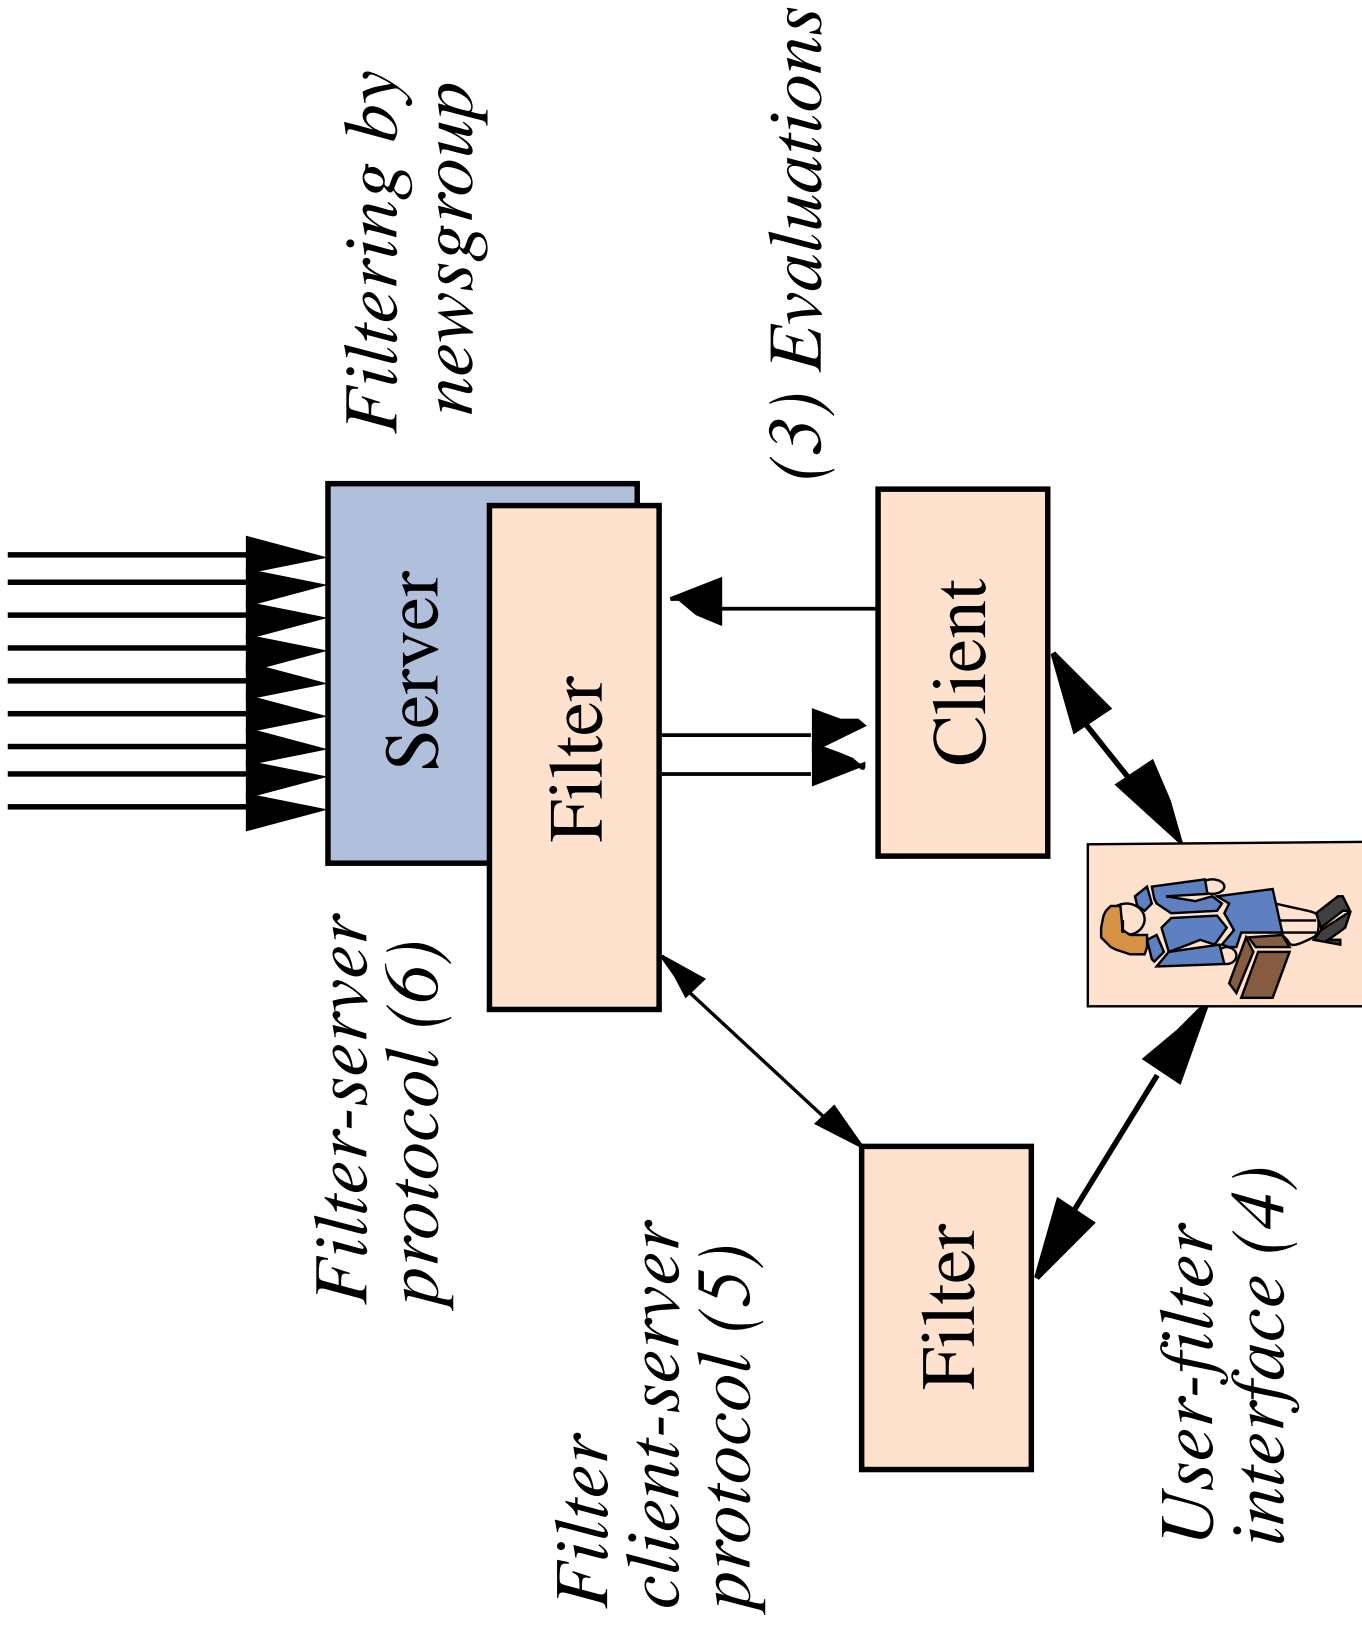


Filtering in the server

- + Can be done in the background.
- + Rejected messages need never be downloaded.
- + Can easily access data bases of documents, other people's ratings, etc.
- User interface more difficult.
- Requires willing service provider.



Possible architecture of filter in the client



Possible
 architec-
 ture of
 filter in the
 server with
 additional
 filter user
 interface in
 the client

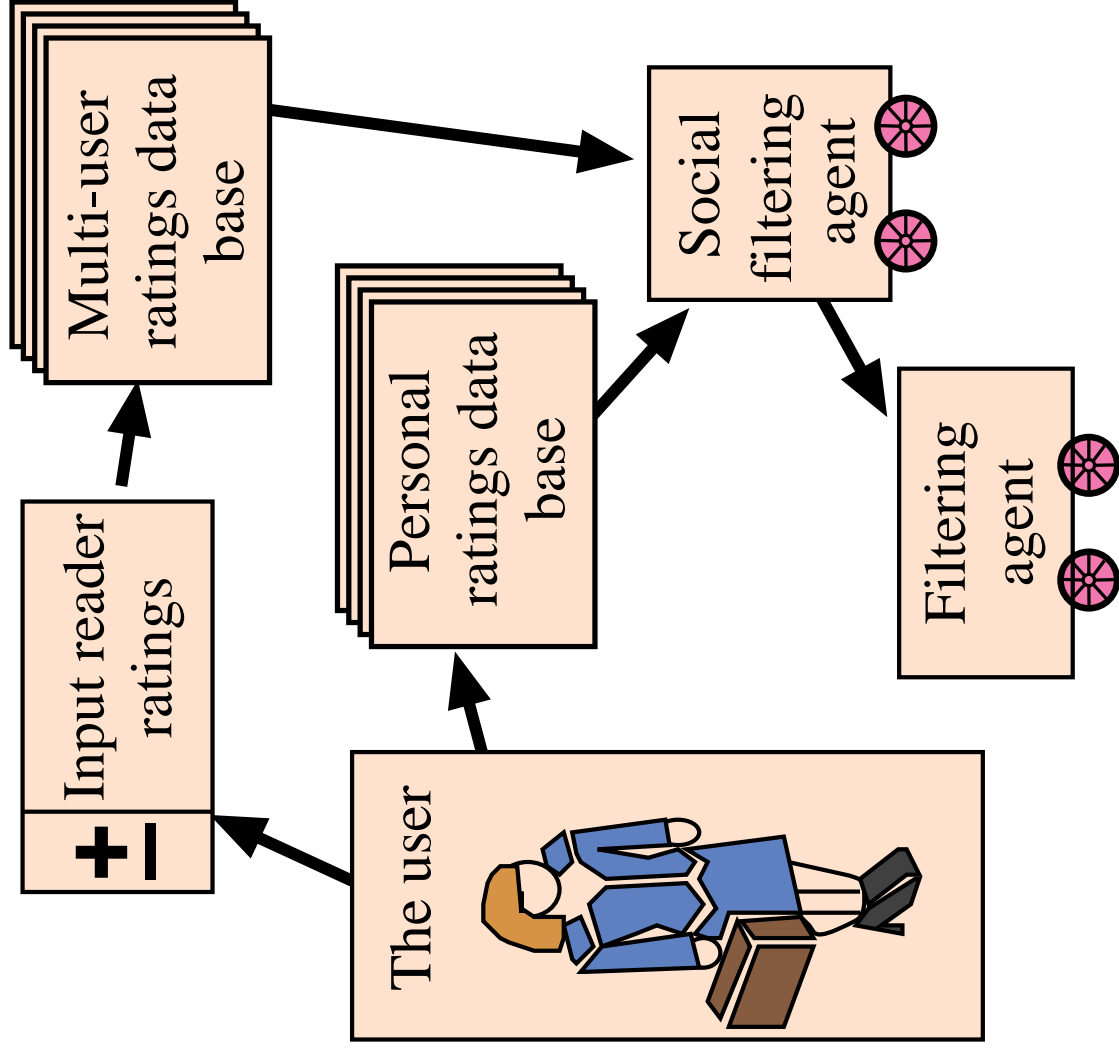
List of new messages in newsgroup A

Checkmark the messages you want to read

<input type="checkbox"/>	Elizabeth Wunderba	18 Nov	Prognosis for Hepatitis B
<input type="checkbox"/>	John Pearson	18 Nov	Re: Prognosis for Hepatitis B
<input type="checkbox"/>	Fred J. Smith	19 Nov	Re: Prognosis for Hepatitis B
<input type="checkbox"/>	Mary von Flüger	08:32	We have a shortage of high den
<input type="checkbox"/>	Fred B. Nelson	09:07	Problem with the X-ray in room

Delivery of filtering results:

- **Sorting of documents in priority order**
- **Sorting of documents by category**
- **Organisation of documents to help manual filtering by the recipient (like the thread lists provided by many news clients).**
- **Rejection of documents, delivery only of positively filtered documents**
- **No sharp limit: Sorting to an almost invisible position is the same as rejection**



Social filtering agent, multi-user ratings data base


Collects people's ratings of documents, store them in a data base.

The filtering agent uses other people's ratings as an aid in filtering.

Issues

Is the social filtering agent part of the normal filtering agent, or a separate module?

How are multi-user ratings used? Average, Median, Upper quartile; for everyone, for those similar to the user or belonging to the same peer group?




User interfaces

- **Input reader ratings.**
- **See and change personal filtering controls and rules.**
- **Retrieve resources using existing resource retrieval systems like mail clients, news clients, Web browsers, Web search engines augmented by filtering and social filtering agents.**
- **Input author ratings.**

How to get users to input ratings

1. Make it very easy to input ratings. Example: Many message systems have a simple keyboard command for “go to the next message”. If pushing any of the keys “0” to “9” could have the same effect, but at the same time rate the last seen document, there will be no extra command to give a rating.
2. You will only get social filtering for your own needs if you also input your ratings.
3. Intelligent filtering will only work for you if you regularly provide your own ratings.
4. (Make people pay for the right of provide ratings, possibly with “virtual money”.)



Intelligent filtering

AI program looks at personal ratings, possibly also at other people's ratings, derive filtering rules, modify the personal filtering settings.

Issues:

- Behind the scenes, or interacting with users?
- Can the user understand the filtering rules?
- Can the user trust filtering rules set up by an automatic agent?

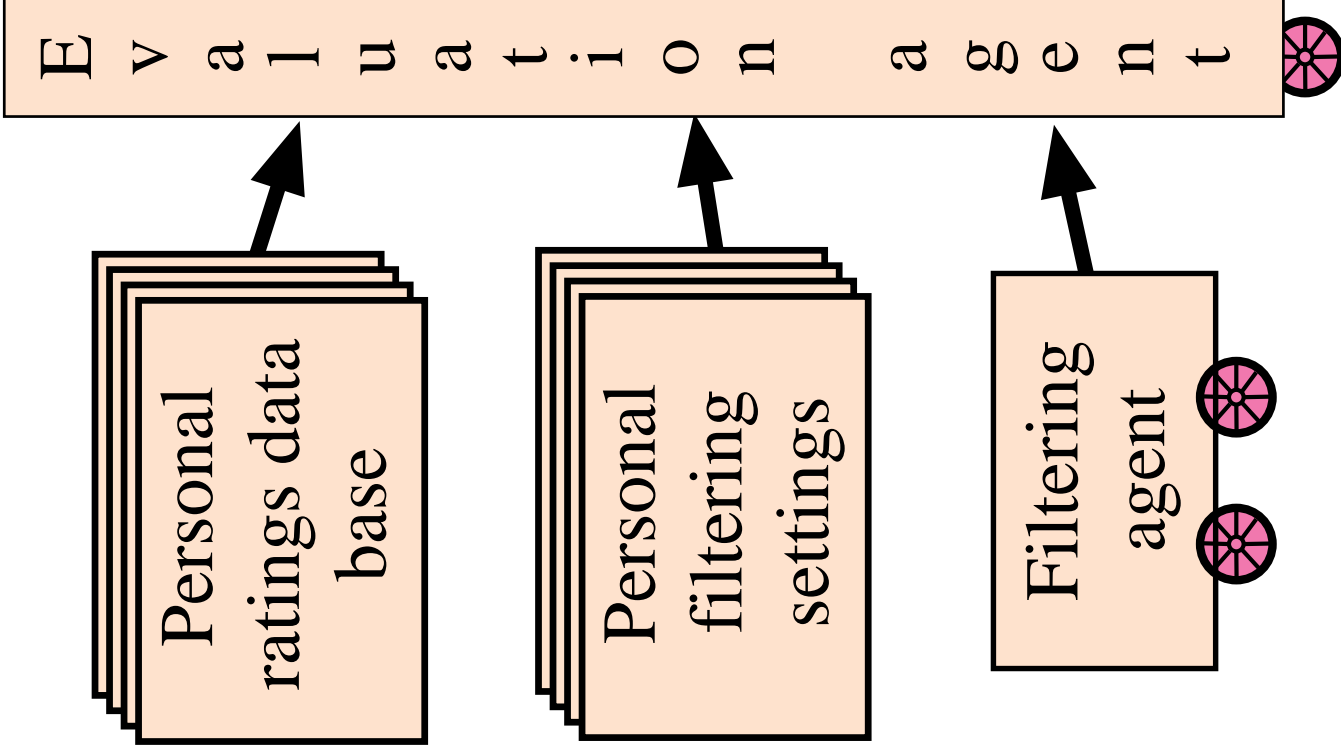
Evaluation

Method 1: Evaluate during actual usage, compare filtering to ratings input by users

Method 2: Collect a data base of rated documents, and test different filtering algorithms on it.

Danger with method 2: Can produce filtering methods which give very good results for the test data base, but are valueless anywhere else, for example a filtering method saying: Give high rates to document no 1, 4, 23, 34, 56, 77 and 83.

This danger is especially true if the filtering rules are derived automatically, as in so-called “intelligent filtering systems” which use AI methods to derive filtering rules from user ratings.



Use of existing standards: Resource delivery to users

Delivery of documents to users is normally done using HTTP and FTP (normal Web lookup), SMTP, POP and IMAP (e-mail delivery) and NNTP (news server-server and client-server protocol). It is an advantage if we use the same protocols as in other Internet usage as much as possible, since this makes it easier for other people to extend their software (news clients, mail clients, Web browsers) to use our filtering services.

If we want to use only one protocol:

IMAP Powerful protocol, can be used to retrieve mail, news and Web pages, has better news control than HTTP, but not widely adopted.

HTTP Users can use ordinary Web browsers, for example using an extension of the Web4Groups Web gateway as delivery tool.

Use of existing standards: Other uses

- HTTP** Can be used for many needs as the protocol to rapidly get and put small units of information. Can be modified, for example change “HTTP/1.0” in the first line to “SELECT/1.0” for our own protocol, and use only a limited set of all facilities HTTP can provide.
- PICS** Can be used to specify rating systems and permitted values. Can also be used as a format to convey ratings. For example: PICS already specifies formats to put PICS labels in the <HEAD> of HTML documents and in the Heading of RFC822 e-mail messages. PICS also specifies formats to get PICS labels from PICS servers, so-called label bureaux.
- The current PICS user interfaces are oriented toward forbidding the user to see forbidden texts, which is not so suitable for general social filtering needs. But PICS is still a good choice if we modify this part of PICS to suit our needs.

Example of a PICS definition of a rating system suitable for collaborative filtering systems:

```
((PICS-version 1.1)
 (rating-system "http://select.org/v1.0/reader-ratings" )
 (rating-service "http://www.select.org/v1.0/")
 (icon "icons/select-icon.gif")
 (name "SELECT Internet Social Filtering Service")
 (category
  (transmit-as "select-reader-rating")
  (name "Select Reader Rating")
  (min 0)
  (max 9)
  (label (name "abstain") (value 0))
  (label (name "totally-value-less") (value 1))
  (label (name "very-uninteresting") (value 2))
  (label (name "uninteresting") (value 3))
  (label (name "not-very-good") (value 4))
  (label (name "of-some-value") (value 5))
  (label (name "better-than-average") (value 6))
  (label (name "interesting") (value 7))
  (label (name "very-interesting") (value 8))
  (label (name "exceptional-superior") (value 9))
 ))
```

Example of a PICS label which a user sends to rate a document:

```
(PICS-1.1 "http://www.select.org/v1.0/user-ratings"
by jpalme@dsv.su.se
labels on "1997.11.01T12:44+0200"
for "http://euroseek.net/"
ratings (select-user-rating 9))
```


Example of a PICS definition of rating system for self-rating of own documents

```
((PICS-version 1.1)
(rating-system "http://select.org/v1.0/author-ratings")
(rating-service "http://www.select.org/v1.0/")
(icon "icons/select-icon.gif")
(name "SELECT Internet Social Filtering Service")
(category (transmit-as "select-author-rating")
(name "Select Author Rating")
(min 0) (max 9)
(label (name "Doctoral-thesis-or-equivalent") (value 9))
(label (name "paper-accepted-in-peer-reviewed-
scientific-journal-or-equivalent") (value 8))
(label (name "abstain") (value 0))
(label (name "masters-thesis-or-equivalent") (value 7))
(label (name "other-scientific-research-paper-or-
university-textbook") (value 6))
(label (name "other-well-researched-and-
peer-reviewed-document") (value 5))
(label (name "official-government-document") (value 4))
(label (name "fiction-or-poem-or-song") (value 3))
(label (name "discussion-item-in-newsgroup-
mailing-list-or-equivalent") (value 2)))
```

```
(label (name "other-document") (value 1)) )
```

Example of using a modified HTTP plus PICS to send in a rating to a multi-user ratings data base

```
PUT-USER-RATING /ratings-input SELECT/1.0
(PICS-1.1 "http://www.select.org/v1.0/user-ratings"
by jpalme@dsv.su.se
labels on "1997.11.01T12:44+0200"
for "http://euroseek.net/"
ratings (select-user-rating 9))
```

Example of a self-rating in the <HEAD> of an HTML document

```
<HTML><HEAD>
<META http-equiv='PICS-Label' content='
(PICS-1.1 "http://www.select.org/v1.0/author-ratings"
by jpalme@dsv.su.se
labels on "1997.11.01T12:44+0200"
for "http://www.dsv.su.se/~jpalme/select/rating-
choices.html"
ratings (select-user-rating 5))'>
<TITLE=...
```

Spamming problem

Spamming	Forcing or cheating people into getting information they do not want.
Mail spamming	Sending unsolicited ads in e-mail, usually with falsified sender information, often misusing mailing lists.
Search engine spamming	Giving a Web document incorrect search key information to cheat search engines into showing this Web document before others.
Social filtering spamming	Cheating social filtering systems into believing that your documents have been highly rated by users.

This is a difficult problem. Possible solutions:

1. Require raters to get an account, do not allow them to use the account until the password or account name has been sent to them by e-mail.
2. Only allow use of ratings for members of special societies.
3. Employ people who are paid to do the rating.
4. Require all raters to get a cryptographic certificate and use this when submitting ratings, but this might make it too difficult to input ratings.

