

# Diffusing Information In Organizational Settings: Learning From Experience

**Dave Snowdon**

Xerox Research Centre Europe  
Grenoble, France  
Dave.Snowdon@xrce.xerox.com

**Antonietta Grasso**

Xerox Research Centre Europe.  
Grenoble, France  
Antonietta.Grasso@xrce.xerox.com

## ABSTRACT

Recommender systems selectively circulate information enriched with comments and feedback based on people's experience. These systems filter information in a semi-automatic and high-quality way in order to support a community during their work or leisure practices. However recommender systems are usually separate tools that require a degree of effort to be used, both when receiving information and to insert new feedback. In this paper we present our informal experiences with the use of multiple user interfaces (interactive large screen, email, paper and PDA) as means to improve the diffusion of information through an organizational unit and to improve access to information stored within an existing recommender system.

## Keywords

Information sharing, serendipity, large-screens, recommender systems

## INTRODUCTION

Early recommender systems, for example Firefly [9] and GroupLens [10] provided personalized recommendations of movies, restaurants, newsgroup articles and music to an Internet audience. These early systems generally used statistical algorithms to perform what is called in the literature automated collaborative filtering. Basically, the algorithms learn to identify groups of people with similar preferences within given domains of interest or genres. As a result, recommender systems are able to provide personalized recommendations, predicting how much an item is likely to appeal to a user based on how others evaluated the item. The more items a user evaluates, the better the system will be able to personalize its recommendations.

Recommender systems based on collaborative filtering have since found a niche in electronic commerce. Many on-line merchants now provide personalized recommendations using such techniques (e.g., Amazon.com), however in the context of this work we are interested in the potential of

using recommender systems for sharing knowledge in an organization [2][3]. In this context, recommender systems provide the following beneficial features:

- Mechanisms to annotate information and quickly share relevancy evaluation within a community.
- Support for both focussed and serendipitous information discovery, enhancing the social network.
- Mechanisms to trace who is active and knowledgeable in specific domains.

As an attempt to improve information sharing in our research centre a web-based recommender system was deployed which aimed to provide users with the above-mentioned benefits. However there was some evidence (mainly from informal interactions and usage monitoring) that for a large number of users the effort to use the tool was still too high. This led to the development of a collection of new user interfaces to receive recommendations and to annotate the information – this collection is known as the CommunityWall due to the principal interface being a large interactive screen.

We now have sufficient experience of this combined system to present some informal evaluations of its effectiveness within a community of about 100 people. We hope that these observations will prove useful to other practitioners, particularly those interested in community building and knowledge sharing.

In the remainder of this paper we will describe the integrated system produced by the composition of the web-based recommender system (Knowledge Pump) and the large screen display (CommunityWall). We will then present our observations on the effectiveness of our approach and compare with current related work.

## DIFFUSING INFORMATION

The resulting overall system (shown in Figure 1) is composed from two prototypes developed here at XRCE. The first, Knowledge Pump (KP) [3] is a web-based shared bookmark and recommender system that uses collaborative filtering to provide recommendations tailored to each user in a certain community. It has been in use since 1998 and currently has 375 registered users spread among several sites – of these less than 100 are active users. Recommendations are sorted by community since people who may agree in one domain may violently disagree in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CHI 2002, April 20-25, 2002, Minneapolis, Minnesota, USA.

Copyright 2001 ACM 1-58113-453-3/02/0004...\$5.00.

another (We used the term “community” instead of simply domain of interest, because we wanted to underline the fact that a community is a set of domains of interest *plus* the people in the organization with that set of interests). KP allows users to submit recommendations of URLs, local files (via upload), or text. A recommendation consists of a rating and, preferably, a comment, along with the user’s classification of the item into one or more communities. In turn, the Knowledge Pump calculates a personalized set of recommendations for a user for each community to which s/he belongs. A selection of recommendations offered via Knowledge Pump is shown in Figure 2.

Each recommended item consists of a link to the item, the predicted score for the user, a list of the names of the users who reviewed it and links to their comments. These last components, reviewer names and comments, are what bring out the potential of a workplace recommender system. A system like Knowledge Pump allows users to track who is active and knowledgeable in their community as well as in turn providing them ways to build their own reputations. Furthermore, the comments around a recommended item enhance the value of that item (in a sense, turning information into knowledge). Lastly, the patterns of recommendations and contents of recommended items can be mined to establish foci of interest and discover who are the local authorities.

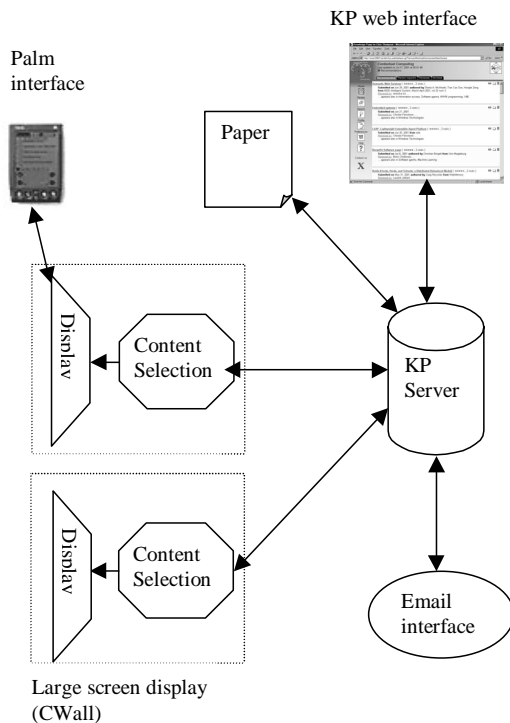


Figure 1: The basic architecture of the system.

When submitting a new item or reviewing an existing one users can record their opinions via numerical ratings representing opinions from “hated it” to “one of my

favourites” and also attach textual comments which explain their rating, suggest related material or justify why they submitted the document. Although the system can be considered to be a success with a relatively large (for an internal prototype) user base, the majority of users at XRCE do not use it. The most commonly cited reason for this is because the interface is simply too time consuming and requires that users navigate to the KP page in order to use it.

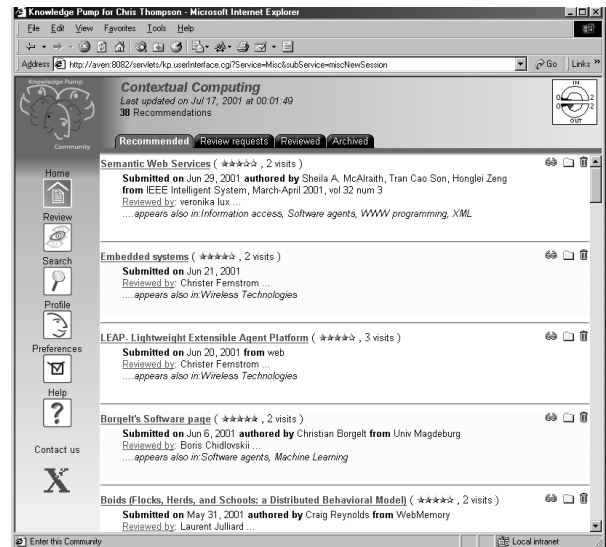


Figure 2: Knowledge Pump (KP) screenshot showing a list of recommended items for a user in a community.

In response to some of these comments and also to experiment with other interfaces we developed the CommunityWall (CWall) large screen display system with the goal of creating a system that provided low effort ways of finding and submitting information and more of a social focus for people on-site. We anticipated that KP (which provides powerful search facilities) would be good when users had a specific information need and we envisaged CWall as supporting a more serendipitous form of information discovery. We used two approaches in order to achieve this: a large screen display for presentation, and multiple input methods.

### Large Screen Display

We chose a large screen display as our main output mechanism in order to have a display that everyone would see in the course of a normal working day. Houde et al [6] also used a large screen as a means to disseminate community related information. They observed that while their large projected display in a communal area was successful, no-one chose to view the same information at their desk via a web-browser. We were hoping that by using a large screen in a communal space we could reach people not using KP. By providing a display large enough for several people to view it at once we also hoped to create a social focus and a mechanism for triggering conversations.

Initially the display was situated next to the black & white and colour printers used by members of our workgroup which also happened to be on the corridor giving access to the office of the lab manager and the group's assistant – we were therefore guaranteed that many people would pass by the display in the course of their working day. After about 10 months of use we moved the display to the entrance foyer of the building used to host visitors to the research centre and later installed a second display in the cafeteria. We found that each site generated different user expectations, which we will discuss in the section on lessons learnt.

### Layout

When designing our display we chose the metaphor of a bulletin board to which anyone can post information. The information can be plain text or web pages that include images (including animated GIFs). We wanted to avoid a standard desktop look & feel and so items are placed randomly on the screen to give a more “organic” feeling. However, in order to increase readability items are not permitted to overlap. The large screen is touch sensitive<sup>1</sup> and people can interact with it using their fingers – the controlling PC's keyboard and mouse is hidden from view. Each item displayed on the screen shows at least a title, a number of stars indicating the average numeric rating, an icon indicating the number of comments on this item, the name of the person who submitted the item and some icons allowing the user to perform specific actions on the item.

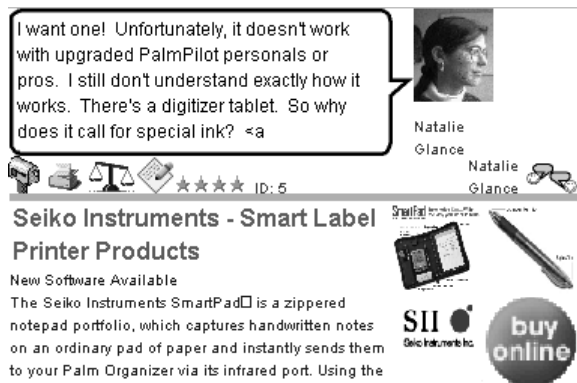


Figure 3: Close up on an item displayed on the CWall showing a comment, the picture and name of the person who submitted the comment, some images from the page, the rating (4 stars) and the first few lines of text.

Items which a CommunityWall instance decides are more relevant (see the section on content selection) at a given moment may also have images selected from the web page displayed, the first few lines of text from the page and a

<sup>1</sup> We currently use a back-projected SMARTBoard for one of display and a plasma screen with a SMART overlay for the other.

display of any comments submitted by other users. An example item representation is shown in Figure 3.

### Interaction

Interaction with the system is deliberately simple, users can:

- Touch the item anywhere except for the four action icons to cause the system to expand the amount of space given to the item (up to the limit of about 30 lines of text) in order to find out more about an item.
- Touch the mailbox icon to email the item to themselves or someone else (a palette of faces is displayed and the user selects the face of the recipient).
- Touch the printer icon to print the document at a nearby printer.
- Touch the balance icon to record a numeric rating (the user is asked to identify themselves by selecting their image).
- Touch the pen and paper icon to write a comment on the screen. We use software delivered with the SMARTBoard to convert the handwriting to text.

Figure 4 shows a screen shot of a typical display. One point to note is that we make no attempt to display the full text of a document on the large screen – we decided at the beginning that a public space was probably not the ideal location to read possibly long documents and we wanted to keep interaction quick and simple and avoid more difficult interaction such as manipulating a scrollbar. Our idea is that we should give people enough information to decide whether an item is interesting and then if they want to read more they can email the item to themselves for reading on their desktop, print it, or capture it on their Palm Pilot<sup>TM</sup> (see below).

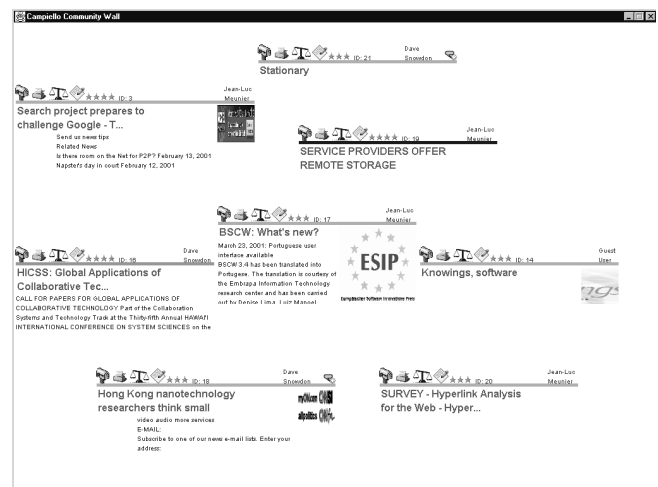


Figure 4: CWall screenshot

### Content Selection

Since only about 10-15 items can fit on the display at one time we implemented a mechanism for automatically selecting items from the database. Items can be classified into a number of types (sticky note, paper, meeting,

conference announcement, etc) and each type can be associated with a rule which is responsible for assigning a priority to items of that type. For example the relevance of a scientific paper might be linked to its numeric rating and the number of comments it has received while a meeting announcement might increase in priority shortly before the time of the meeting and then fall to zero afterwards. We implemented a number of simple rules and the possibility to compose them hierarchically so that administrators of a display can compose composite rules using a simple text configuration format without having to understand how to program. At intervals (currently every 10 minutes) the system reapplies the rules to any active items (those that have been created, rated, commented on or interacted with in the recent past) and selects those with the highest priority for display. This means that the display changes often enough that people who pass by several times in a day should see different items. It also allows us to give priorities to different types of item at different times of day.

Our current rule-set includes a rule that adds a small amount of random noise and a rule that decreases the priority of items according to how much time they have already spent on the display in order to prevent a few items from monopolising the display and to ensure that users see a variety of items if they pass by the display several times in one day.

### Multiple Input Methods

We formed the hypothesis that one reason why KP was felt to be too time consuming was possibly because the web interface was not always a good match for the medium carrying the information which people wanted to submit. In order to reduce the cost of submitting information we have therefore developed several additional user interfaces that can be used to submit information to the system.

- Email : Many hi-tech organisations have a strong email culture; our centre is no exception and so like Houde et al. [6] we decided to allow people to submit new items by email. This means that documents such as calls for papers which are often received by email can be forwarded to the CWall<sup>2</sup> using any mail program. Also the CWall can be put in copy of messages sent to others.
- Web bookmarklets: In order not to disrupt browsing, forcing users to go to a particular page in order to submit an item, we created bookmarklets<sup>3</sup> that would capture the current page URL and title, and pop up a window allowing the user to confirm the selection and optionally add more information such as a numeric

<sup>2</sup> Keywords can be added to the email to give a comment, rating and indicate an expiration date.

<sup>3</sup> A bookmarklet is a small JavaScript program stored where the URL would be in a normal web browser bookmark. People can use these in exactly the same way as a normal bookmark.

rating or comment. Cookies are used to prevent the user from needing to login after the first time the system is used.

Figure 5: Paper form allowing users to submit hardcopy documents to the CWall and KP. The form allows users to check boxes corresponding to type of document, a rating and their telephone extension. They can also leave hand-written comments in the box in the middle - these are displayed as images on the CWall and in KP.

- Paper: Using Xerox's Flowport™ software we constructed paper forms that allow users to submit hardcopy documents. Using a combined printer/scanner/copier next to the large screen<sup>4</sup> users are able to place a cover sheet (Figure 5) on top of the document and press the start button on the copier. The cover sheet is recognised as introducing a new document, the document is scanned, OCR'd and then the keywords in the OCR'd document are analysed to decide which KP community it is best suited to. Users have the option to override the automatic classification using the normal KP web interface.
- Palm Pilot™: One current disadvantage with the CWall is that touch sensitive back projected or plasma displays are still expensive and so we wondered if we could have a useful system using a non-interactive projected display such those found in many meeting rooms. Since many people now have PDAs such as the Palm Pilot™ we thought that these could become a sort of personal interface to the system. We prototyped a Palm application that uses infrared communication to retrieve a list of items currently on display on the screen. Users can read the items, comments and ratings on the items submitted by other users, rate and comment on items, print the items and email the items

<sup>4</sup> This is the main printer/copier for our workgroup so we were confident that people would come here anyway with paper documents to copy or collect from the printer.

to themselves or others. We also allow Palm users to create new items that could be submitted to the CWall via infrared. The Palm application (shown in Figure 6) has an on-line mode for when the user is current in front of the display and an offline mode which allows them to browse existing items and create new ones for later transmission to the CWall.



Figure 6: Palm interface to the CWall

- KnowledgePump: Although initially conceived of as a separate system we did not want users to have to choose between submitting documents to the CWall and KP so we established a link between the two systems such that documents submitted to one would be available to the other.

### LESSONS LEARNT

Firstly, we learnt that the effort required to prevent people from using a system is surprisingly low. Even though people are accustomed to sharing information and papers by email there is some resistance to using software such as Knowledge Pump. Even though a core group of users has found it of benefit others say that the effort required to go to the web page and open a form for submitting a document or reading their recommendations is enough to put them off using it. These users will typically however make great use of email to recommend and exchange information - this is what led us to create an email interface for the CWall.

After some initial scepticism the CWall received a favourable response from most users and we received several suggestions for improvement. We currently have two CWall displays running at our site, one in the cafeteria and one in an entrance foyer. On an average day two documents are printed and users send two more to themselves by email. We estimate that we have 20 regular users and a few other occasional users. After the system had been in use for about a year we performed an analysis of the log files and asked both users and non-users to fill in a questionnaire designed to understand why people did and did not use the system. This section will present what we learned from this and from informal conversations with users.

From the questionnaire it emerged that most of the users who claimed to use the CWall claimed not to use the

Knowledge Pump<sup>5</sup>. This was reassuring since it seems to mean that the CWall fulfils a real need and distributes information stored in KP to people who might not otherwise see it. Likewise documents entered via the CWall can become available to KP users. From log analysis we also discovered that there is a correlation between the number of items submitted to the system in a day and the number of items examined on the large screen display. We interpret this to mean that users are aware of and react to new documents appearing in the system.

We also found that, unsurprisingly, people want to be sure that their submission has been seen by others - there is an element of public performance, having submitted a document it is necessary that it finds an audience. After submitting something to the CWall users would pass by the screen in order to see if it was displayed. However, because the rules didn't guarantee that new items were displayed immediately users might not see them then. We found users less motivated to submit items if they could not see them on the screen, because then they did not know if other people had seen them. To correct this we added a new sub-rule to the rules for general documents that boosted the priority for new items to ensure that they would get seen at least for a while irrespective of their other merits. We are also considering other ways of giving feedback such as sending users reports of how many other people have visited documents they submitted.

### User Categories

Results from the questionnaire led us to think in terms of three categories of people:

1. Active users – these are willing to invest time in prototype systems and to invest some effort in sharing information with others. These people are typically users of the KP web-interface.
2. Interested users – these are interested in information system but don't want to invest much personal effort. This group will typically not use the KP web-interface but will use email (and the CWall email interface) to share information and the large screen to find interesting information.
3. Inactive “users” – these don't want to use some information sharing tools. Typical reasons given for this was a reluctance to use tools that could potentially gather information about their habits or because they did not feel part of a community and did not feel that they would get anything in return for their efforts. Here is a typical comment from this category of user “*I am*

<sup>5</sup> In fact they were using Knowledge Pump indirectly since this is where a large proportion of CWall documents come from. However, they did not need a Knowledge Pump login or to be aware that they were using information from Knowledge Pump since they saw only the CWall large screen and (and possibly email) interface.

*so proud that I hate the idea that somebody or a system can profile me, classify me in communities, pair me with people... also I dislike that this information can be used ultimately to push me things, sell me things (or even worse) etc."*

### **The Effects of Location on Users Expectations**

One unexpected result was that users' expectations of the system appear to change according to the location of the display even when location changes are restricted to one workplace. When we installed the first system in proximity to our offices and those of our workgroup it was our workgroup that mostly used the system. There was a small but loyal group of users with most submissions being work related but also some being jokes or some sort of fun item. Comments received were pertinent. After about ten months we moved the display to the foyer of a building which is used to host visitors as well as office space for about 40% of the people on site. This move was done in order to facilitate giving demonstrations to larger groups of people and also to increase the number of potential users by placing the system in a place with more traffic. We found that people were more likely to be curious and experiment with the system. But also that because of the change of building we lost some of our previous user group. At the same time we got more "junk" comments submitted from the screen as people would experiment with the both the system and the handwriting recognition - for the first time we had to contemplate cleaning the database to remove these spurious comments. We also hypothesise that because of the change of location there was less sense of ownership of the system as it was no longer physically associated with the group which had created it. Three months after the move we installed a second screen in the site cafeteria hoping to capitalise on the large numbers of people who use the room during the day for eating lunch, tea breaks and informal meetings. Although this did appear to attract new users we also found that they seemed to have different expectations of the information displayed on the screen. Whereas the first two locations had been associated with work, the cafeteria was associated more with leisure and so we found users requesting more leisure items be displayed - such as cinema timetables and reviews of current films. In order to respond to this demand we are contemplating mining local cinema timetables from the web in order to automatically generate this information. Since each display has its own rule set for content selection we can already change the rules for each display so that the one in the cafeteria is more predisposed towards leisure while the others are more predisposed towards work - we are in the process of identifying appropriate rules.

### **Identification and Trust**

Our initial identification mechanism worked by asking users to select their picture from a palette of users whenever they performed an operation that needed their ID (such as rating an item). This worked well initially within our workgroup (membership 15-20 people) where we trusted people to be responsible and everyone knew each

other well. An unplanned positive side effect of this was that people could recommend items to others by choosing the email icon and selecting another person's face. However, when the system was moved into the entrance foyer the number of possible users expanded (about 100 potential users) and made this mechanism cumbersome. Also "junk" comments started to appear indicating that trusting people to be responsible was no longer working. We are now investigating the possibility of using a badge reader so that the identification badges used to control access to the building also identify users to the CWall.

The Palm application is working but since infrared only allows one Palm to connect at a time we are currently experimenting with Bluetooth<sup>6</sup> to allow up to seven users to connect to a display at once.

### **RELATED WORK**

There are several other groups developing systems based on large displays, here we will describe the most relevant. Not all systems have the same goal as ours and we can categorise known systems into three basic groups

- Those that support synchronous collaborative work, which is normally co-located: Dynawall [1], Tivoli [8].
- Those that support more loosely coupled collaboration but have a bias towards synchronous collaboration, though not necessarily co-located: Notification Collage [4].
- Those that are intended for purely asynchronous operation: Apple Newspaper [6], CWall.

Of the systems we know about the Apple Newspaper and the Notification Collage have the most in common with the CWall and we will briefly compare the three systems in more detail.

#### **The Apple Newspaper**

Houde et al. wanted to improve communication among a group of researchers at Apple research. Realizing that they could not improve communication through technological means alone if the social conditions were not appropriate they started a tradition of getting together in a communal lounge at tea-time for tea, coffee and cookies thereby providing group members with an informal opportunity to get together and talk. To complement this they created a space on the lounge wall where people could post information about their projects. These pages could have stickers with barcodes attached that would link to a web page with more information. After some initial enthusiasm use of the wall quickly fell off as it was felt too much effort to produce good-looking web pages. In order to address this problem they created a system that would accept email submissions and automatically format them into a web page

---

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

that represented the front page of an on-line newspaper. This front page was projected onto a wall in the lounge and community members could also view the newspaper at their desktops using the web. This second system was more successful and the use of email was a good match for the prevalent "email culture" in the research group. However, they found that all use of the newspaper was via the projected display in the lounge area and no one made the effort to view the web page at their desktop. Since the Apple system had a non-interactive display only the full text of one item could be displayed at a time and a new item would displace (for good) whatever was currently on display. In contrast the CWall acts more like a bulletin board and can display summaries of 10-15 items at once – users can interact with the display to find out more. Like the Apple newspaper we allow people to make use of email to submit items to the system and in addition we also allow people to receive items via email.

### **The Notification Collage**

Like the CWall, the Notification Collage (NC) [4] uses the metaphor of a bulletin board which can display several randomly placed items simultaneously. The NC can handle several types of media including series of still images, near real-time video, web pages, sticky notes and desktop thumbnails. Unlike the CWall the NC allows items to overlap, with newer items always going on top. However, users can reposition items on the display. In addition a portion of the NC display is reserved for manual placement of items and users can manually place items that they do not want to be obscured in this area. Greenberg and Rounding report that this reserved area is mainly used for video windows of other users so that they can see if people are present and get a feeling for their current level of activity. Another difference between the CWall and NC is that while CWall items may contain animated items they are essentially static - if a web page is updated the new page will be displayed by the CWall when it next fetches the page but there is the assumption that items change relatively infrequently (e.g. less than once per day). On the other hand the NC allows items to be updated as frequently as users desire and users can use the sticky chat feature to have public chat sessions. Finally, the NC relies on custom client to submit items to the system whereas the CWall requires only an email reader and/or a web browser and does not require users to navigate to a particular web page in order to submit new items. Greenberg and Rounding report that video was the most important media type and that users would make heavy use of the video items to check on people's presence and availability. The CWall was never intended to support this sort of synchronous use. In contrast to the Apple Newspaper Greenberg and Rounding report that NC was used at the desktop, but this seems to be limited to users who have two monitors allowing them to dedicate a monitor to NC while continuing to work on the other.

### **Scalability**

When considering such systems for community awareness one important issue to consider is the scalability of the system. Scalability refers to how gracefully performance degrades as load increases and can be considered in terms of the number of users supported, the number of displays that can be served by the system and the amount of total information managed. As the number of users increases the system needs to be able to handle all information submitted by each user and provide enough information to satisfy as many users as possible. If multiple displays are required then more strain may be placed on any central server components and the network infrastructure between displays (clients) and server – these infrastructure limits may prevent the system supporting more than a particular number of concurrently active displays. Finally, as the amount of information increases the system has to work harder in order to present information that is relevant and display a suitable volume of information such that most if not all of the most interesting items can be seen (though perhaps not simultaneously). We will now look at the Apple Newspaper, Notification Collage and the CWall in order to see how they might scale in each of these three dimensions. Without access to the other systems we cannot make quantitative performance measurements but we will attempt to make qualitative estimations of how each system might perform.

Houde et al. do not mention the size of the community supported by the Newspaper but the obvious limiting factor was that only one item could be displayed at once. This requires that either the community is small enough that not many items are generated per day or largely inactive with most people being passive consumers of information. However, since the client is simply a web browser the system should be able to support hundreds of displays (projected or otherwise).

Greenberg and Rounding state that their prototype has been used only within their research group but neglect to give the size of the workgroup. From the seating plans and other information about remote users given in the paper we estimate that the size of their user community was about 8-12 users. As the contents of items changes their directory server needs to notify clients of updates. However, given the experiences of the Collaborative VR community which has much more stringent update and scalability requirements [5] this is unlikely to be a limiting factor unless the number of displays increases beyond several 10s of displays and very heavy use is made of video – in fact the bandwidth required to transmit video is likely to be a problem long before the load on the central directory server becomes problematic. In terms of users and the amount of information the NC's main limit is that items are placed randomly and that new items always obscure older ones – this means that as the amount of information increases each piece of information will be visible for a smaller amount of time. Given that the use of the NC seems to be biased towards dynamically updated documents and synchronous

awareness this may not be a problem in practice as long as the display can accommodate enough video windows and sticky chats to meet the needs of the currently active users.

CWall falls somewhere between the Apple Newspaper and NC in terms of architectural scalability and is probably the most scalable in regarding the amount of information that can be handled. CWall does not really support dynamically updated items and so places a relatively light load on the central component that manages the database. On the other hand because of tracking user actions such as comments and ratings and because each display can handle 10-15 items at once it is more computationally demanding than the Apple Newspaper.

In terms of the amount of information handled, the CWall should be able to outperform both the Apple Newspaper and the NC although unlike the other systems it does not guarantee that an item will be displayed at a particular time – changing the prioritisation rules can, however, ensure that the system will almost certainly display an item soon after its creation or a some fairness metric could be implemented to ensure that popular items receive a fair amount of display time. We already use a rule that decreases the priority of items proportional to the amount of time they have spent on screen relative to other items. By using measures from users such as the number of comments on an item and the average numeric rating the CWall can prioritise documents that the user community is likely to find interesting. Also since each display can have its own rule set displays can be customised to their location.

### CONCLUSION

In this paper we have presented a system designed to diffuse information amongst a community of researchers in which the effectiveness of work relies of good awareness of domain activities and the actions of peers. We started with an existing web-based recommender and added large screen, paper, email and PDA interfaces in order to make the information more accessible to users at a lower effort. We then described some of the lessons we have learnt after more than a year of the prototype being in use. Finally, we compared our system, the CWall, with related work in this area.

### ACKNOWLEDGMENTS

The CWall was originally developed as part of the European i3 framework Campiello. The current CWall user interface is based on a design from the Domus Academy in Milan. Thanks are also due to Cyrille Roy for work on the implementation, Alain Karsenty for installing the system at XRCE Cambridge and the many people at XRCE who have used and given feedback on the system.

### REFERENCES

1. Geissler, J. Shuffle, throw or take it! Working efficiently with an interactive wall. Proceedings of CHI'98 summary, (Los Angeles USA, April 1998).
2. Gance, N., Arregui, D., Dardenne, M. (1998), Knowledge Pump: Supporting the Flow and Use of Knowledge in Networked Organisations, Information Technology for Knowledge Management, Eds. Borghoff and Pareschi, Springer-Verlag, 1998.
3. Gance, N.S., Arregui, D., Dardenne, M. Making Recommender Systems Work for Organizations. Proceedings of PAAM'99 (London UK, April 1999).
4. Greenberg, S. and Rounding, M. (2000) The Notification Collage: Posting Information to Public and Personal Displays. Yellow Series Report 2000-667-19, Department of Computer Science, University of Calgary, Alberta, Canada. <http://www.cpsc.ucalgary.ca/group/lab/papers/index.htm>
5. Greenhalgh, C. M. and Benford, S. D., Supporting Rich And Dynamic Communication In Large Scale Collaborative Virtual Environments, Presence: Teleoperators and Virtual Environments, Vol. 8, No. 1, February 1999, pp. 14-35, MIT Press.
6. Houde, S., Bellamy, R., Leahy, L. (1998), In search of design principles for Tools and Practices to Support Communication within a learning Community, SIGCHI Bulletin, 30(2), April 1998. <http://www.acm.org/sigchi/bulletin/1998.2/houde.html>
7. Huxor, A. (1998), An Active Worlds interface to BSCW, to enhance Chance Encounters, CVE 1998, June 1998, Manchester, UK, Snowdon & Churchill eds., 87-93
8. Moran, T.P., van Melle, W., Chiu, P. Tailorable Domain Objects as Meeting Tools for an Electronic Whiteboard. Proceedings of CSCW'98 (Seattle USA, November 1998), 295-304
9. Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P. and Riedl, J. "GroupLens: An open architecture for collaborative filtering of netnews," in *Proceedings of CSCW'94* (Chapel Hill, NC, October 1994), ACM Press, 175-186.
10. Shardanand, U. and Maes, P. "Social information filtering: Algorithms for automating word of mouth," in *Proceedings of CHI'95* (Denver CO, May 1995), ACM Press, 210-217.