Department for Computer and Systems Sciences

Stockholm University
Kungliga Tekniska Högskolan

SuperKOM – Design considerations for a distributed, highly structured computer conferencing system

By Jacob Palme and Torgny Tholerus

Published in Computer Communications
vol. 15, no. 8, october 1992 pp 509-518

DSV report no 91-010

Abstract: The SuperKOM computer conferencing system was designed to incorporate many ideas from older messaging and conference systems in a simple and general-purpose internal structure. This paper describes the design goals, and then describes how the internal structure of SuperKOM was designed to meet these goals. In particular, SuperKOM provides a distributed, highly structured conference system and integrates electronic mail, computer conferencing and text data bases in one system.

Keywords: Electronic Mail, Message Handling Systems, MHS, Computer conferencing, Bulletin Board systems, Usenet News, Group Communication, KOM, PortaCOM, SuperKOM, Computer Mediated Communication, CMC, Computer-Supported Cooperative Work, CSCW.

Main author's personal address: Skeppargatan 73, S-115 30 Stockholm, Sweden. Phone: +46-8-16 16 67. Internet mail: jpalme@dsv.su.se.

University address: Department of Computer and Systems Sciences, Stockholm University and Kungliga Tekniska Högskolan, Electrum 230, S-164 40 Kista, Sweden. Phone: +46-8-16 20 00.

Gopher server: gopher://dsv.su.se

World Wide Web server: http://www.dsv.su.se

Table of contents

1. Computer Conferencing, Group Communication, Hypertext

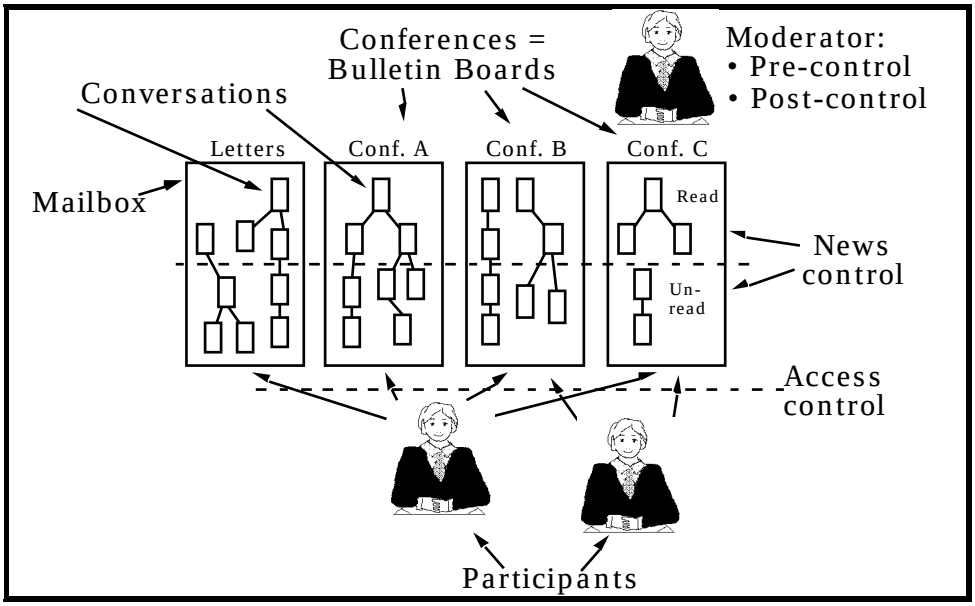

Figure 1
Terminology of computer conferencing

Many different terms have been applied to computer systems for messaging in groups: Asynchronous Computer Conferencing, Group Communication, Bulletin Boards, Hypertext [15], Computer Mediated Communication Systems etc. Most of these systems have the same basic principle in common.

This basic principle (see figure 1) is that communication via messages in a group of people is aided by computer systems which structure the communication space. Common structuring tools are [12]:

Messages or *contributions* which in addition to text may contain diagrams, pictures voice etc. (Multi-media messages.)

Conferences (also called topics, activities, bulletin boards, questions, newsgroups) comprise a set of people and a set of messages being exchanged between these people. The word *conference* will be used in the rest of this paper.

Moderators are one or more people with special control of a conference. In *pre-moderated* conferences, all messages have to be approved by the moderator before being submitted to the conference, in *post-moderated* conferences, any participant can submit messages, but the moderator can remove them afterwards, or move them to a more suitable conference.

Comment links (also called reply links, reference links) connect two messages together where one of them is a reply to, a comment on or in some way references a previous message. A number of messages joined directly or indirectly by such links are called *conversations* or comment trees. The words *comment links* and *conversations* will be used in the rest of this paper.

Keywords or names can be assigned to users, conferences and messages. Users can look up the keywords in *directories* or *indexes*.

Access control are tools to control who may make themselves a member of a conference, who may read and write in it etc.

Hypertext is a system of objects which are linked by defined links. Computer conferencing can be seen as a special case of hypertext, using links to join messages, participants and moderators to conferences and comment links to join messages.

News control is the facility to help a user find which messages are new and unread for that user.

2. SuperKOM design goals

Part of the history of computer conferencing is shown in figure 2, which also shows how SuperKOM was designed with knowledge of many existing conference systems like EIES [6], KOM [9], Parti(cipate)[5] and Usenet News [4].

The worlds first conference systems were Emisari (designed by Murray Turoff) and Forum-Planet (designed by Jaques Vallee) [18]. The idea originally came from the Delphi method used in futurology. Around 1980, a number of different conference systems were designed. Typical of these systems was the centralized view, assuming that all users connect to one central host.

The newest systems (Usenet News, Eies 2 [19], SuperKOM) differ from these systems in their distributed [14] nature, where users are connected to many different hosts, working together via computer networks.

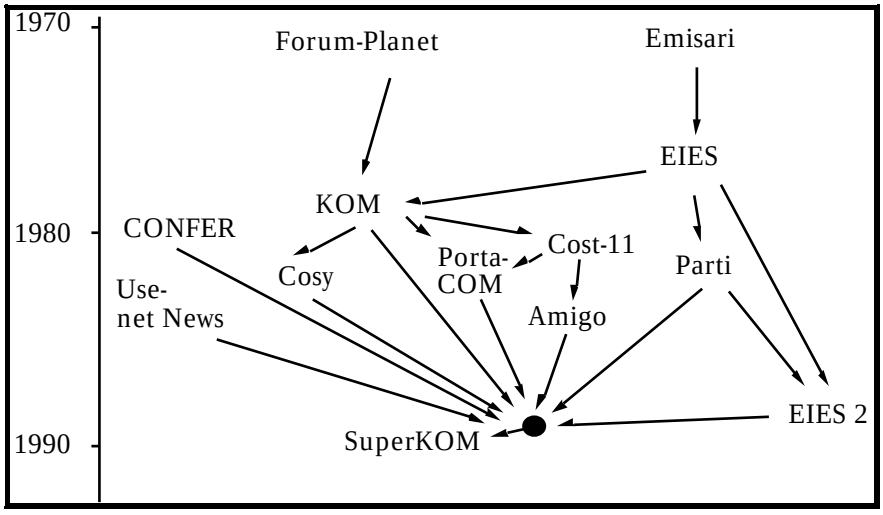

Figure 2
SuperKOM dependence on previous systems

Our design goals were to incorporate in a simple and general-purpose structure the main concepts of several older systems (see acknowledgements chapter below). In addition to the features described in section 1 above, we also wanted to support the following ideas:

2.1. Allow the sending of messages to more than one conference

One and the same message can be submitted or re-submitted to more than one conference. Users should still not be shown the same message more than once, even if they are members of both conferences.

2.2. Integrate personal mail with group communication

Each person has a personal mailbox, which looks very much like a conference with only one member. Thus, a message can be sent or resent to any number of personal mailboxes and/or conferences, and a user would still not be shown the same message as new more than once.

2.3. Natural default recipients for replies

Comments (replies) to messages are by default sent to all recipients of the commented message. The originator of the comment should of course be able to change this default set of recipients. For example, if a message was sent to conference A, conference B and as personal mail to a users C and D, then the default set of recipients for a comment on that message would also be conferences A and B and users C and D. It should also of course be easy to write a reply which is only sent to the author of the commented entry.

2.4. Easy forwarding of conversations to additional recipients

Suppose that in this way, a series of messages are exchanged between the same users. And suppose that one of them decides that some outsider or group of outsiders should take part in this discussion. It should then be simple to send this discussion to additional recipients, which may be users or conferences. In SuperKOM, we chose to give the person, who forwards a message to additional recipients, two options on how the forwarding should work. With one of the options, only one single message and nothing more is forwarded. This option is in SuperKOM called *single copy*. The other option is to forward the whole conversation, or a whole branch of the conversation, to additional recipients. This option is called *for information*. With the second option, also forthcoming, not yet written messages in the forwarded conversation or branch will automatically be forwarded to the additional recipients.

Our experience with SuperKOM has shown that this is a very valuable facility. There often occurs cases, where you want to send whole conversations or branches in them to additional recipients, and to be able to do this with just one single command is very useful.

2.5. Users should be able to join and withdraw from conversations

There should be strong support for conversations. In principle, all which can be done to a conference, should also be allowed for any conversation or any branch of a conversation. Users should thus be able to browse conversations, read news one conversation at a time and be able to become members of (be shown) messages from any branch of any conversation, and to withdraw from (not be shown) the messages from any branch of any conversation. This should, just like for conferences, apply also to not yet written, forthcoming messages in the conversation.

2.6. Hierarchical conference data base

Another design goal for SuperKOM was that the data base of the existing conferences should itself be structured, so that conferences can be organized into groups, departments and other hierarchies of conferences to suit different needs (see figure 3). One subordinate node in the structure can be superior to more than one superior node. Our experience is that this is a valuable feature, but that one should not overdo it, creating an unnecessarily complex branching structure in practical usage of the system.

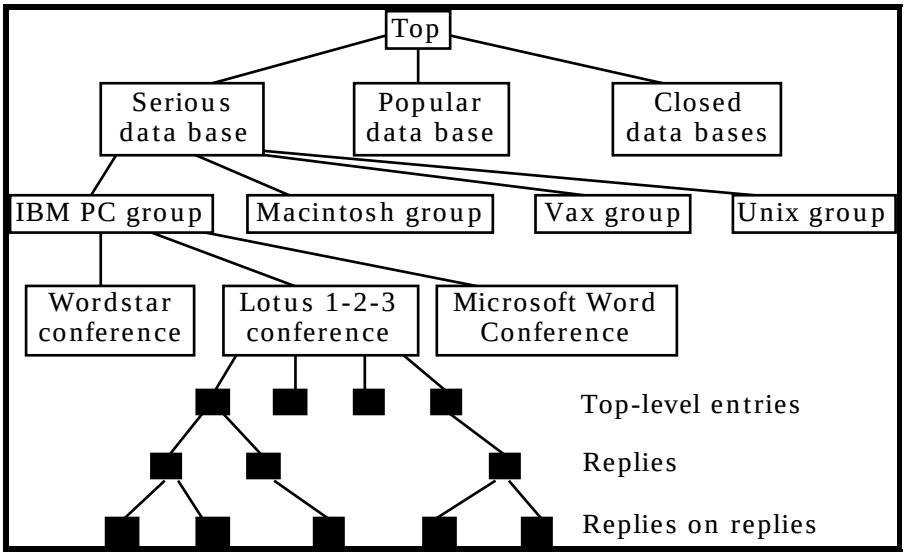

Figure 3
The hierarchical conference data base

2.7. Built-in general purpose directory and information retrieval system

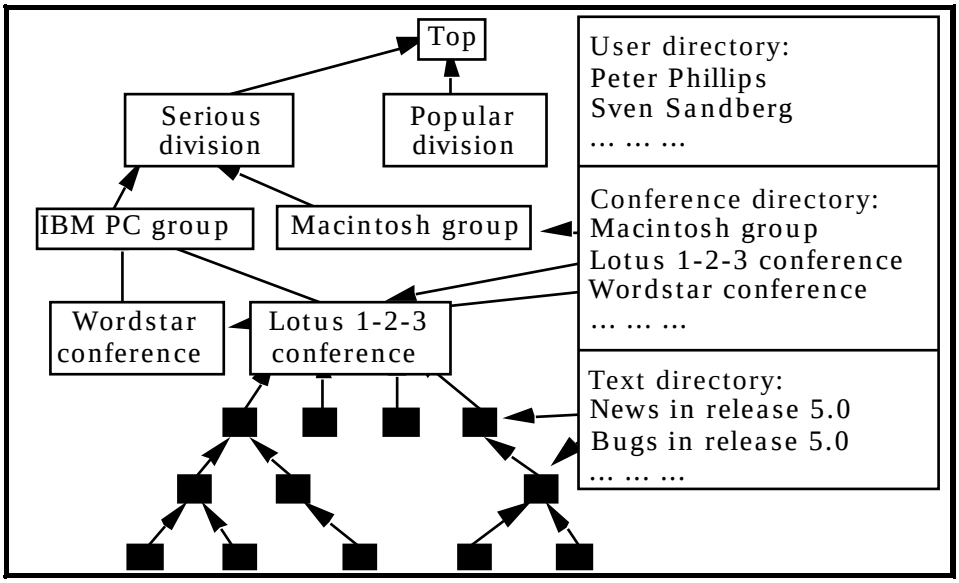

Figure 4
The central directory

It should be possible (see figure 4) for any node in the data base (user, conference, conversation, branch, message) to be assigned a name or keywords in the directory system, so that users can find them by information-retrieval queries.

2.8. A distributed system

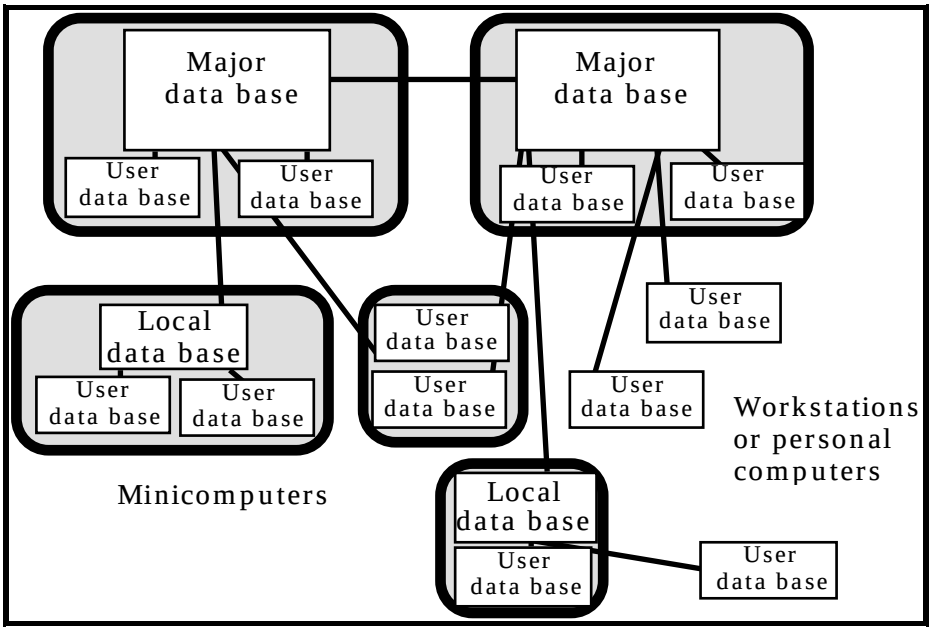

Figure 5
The distributed data base

Finally, SuperKOM was to be distributed (see figure 5). By this we meant that the data base could be distributed on many hosts, each containing only the activities of interest to the users at that host. All features of the conference system should be available also in the distributed environment. For example, if a moderator moves a message from one conference to another more suitable conference, the effect of this action should be applied automatically to all hosts where this message occurs. In principle, the distribution of the data base should be transparent to the users.

2.9. X.400 and and Internet mail compatibility

SuperKOM should allow connections to other mail systems, using the X.400, Unix Mail, Internet and Usenet News messaging protocols. Users in such other message systems are in SuperKOM called *external* users.

It should not only be possible to send mail to external users. It should also be possible to include their names as members of conferences in SuperKOM. If their message system only has electronic mail, and not conferencing facilities, the SuperKOM conference will then appear to them as a distribution list. It should also be possible to nest SuperKOM conferences with distribution lists. The conference would still look as a conference to SuperKOM users, and as a distribution list to external users outside SuperKOM.

It should be possible, from other mail systems, to send messages to SuperKOM, both to individual SuperKOM users and to conferences.

All the facilities described above for handling multi-recipient and multi-group messages, forwarding of conversations etc. should of course also be available for recipients in other messaging systems connected via X.400 or other mail protocols. The

only limitation would be that external users can only use the commands available in their message system. Thus, external users cannot themselves give commands for example to withdraw from conversations.

2.10. A general-purpose and easily extensible internal data base structure

The data base should be structured in a simple and general-purpose way, so that it easily can be extended with new facilities, like for example support for voting, joint editing and other more advanced group communication tasks.

2.11. A modern and fast user interface

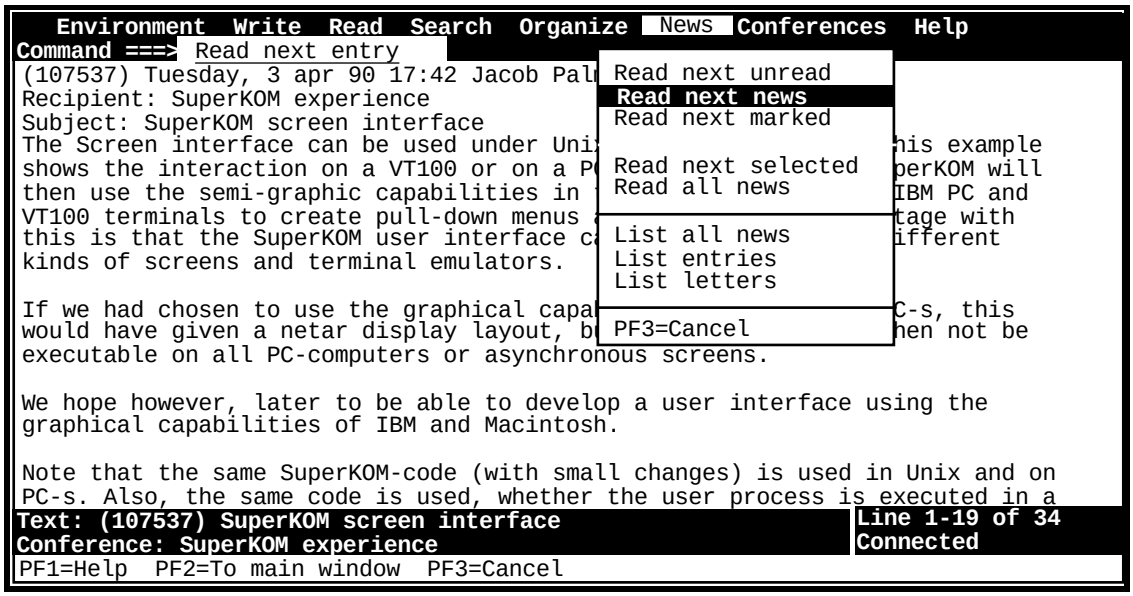

Figure 6
The user interface

Finally, we wanted to give SuperKOM a modern, Macintosh/Windows-like user interface with pull-down menus (see figure 6). We chose to follow mostly the IBM CUA standard for such user interfaces. A line-oriented user interface is also available as an alternative for those users who prefer it.

And we wanted the user interface to give fast response times also when the system is used by many simultaneous users.

3. Meeting the design goals

3.1. System architecture

In order to meet the design goals of a modern screen-oriented and fast user interface, we decided to give every user his own mailbox data base. This is different from most other conference systems, which have only a central data base, common for many users. With a separate mailbox data base for each user, most user commands can be fullfilled via accesses only to the mailbox data base of that user. Access conflicts on the central data base, when used by many simultaneous users, would thus not reduce response times.

A separate mailbox data base for each user can be stored locally on the hard disk of the personal computer or workstation for the user. This will give even faster response times and allow the users to have the kind of modern screen-oriented user interface which they are accustomed to from other modern personal computer software.

Another advantage with a separate mailbox data base for each user is that news control (enabling users to browse only new messages, and ensuring that they will then not be shown the same message more than once) is easier to provide combined with any browsing strategy preferred by the users. Many other conference systems provide news control only if the user reads messages in a consequence in a linear order.

Also, a separate mailbox data base allows the users to reorganize their personal data base into folders according to their wishes. In SuperKOM, folders and conferences are two aspects of the same facility. Conferences can be seen as global folders, shared by many users, but users can also create local conferences, only in their mailbox data base, which then will act as personal folders.

To avoid waste of storage, it is possible to set up the user data base to contain virtual link to large messages, which can then be fetched at read time from the central data base. This is suitable if there are fast direct connections between the computers handling the user and central data base. In the opposite cases, the full text of each message is copied to the user data base.

The program structure chosen is shown in figure 7.

This organization of the system into many processes will give faster response times for most user commands. In many other conference systems, the user is kept waiting while the central data base is updated to enter a new message. In SuperKOM, the new message is just put in a queue, and the central process will update the central database in the background, while the user is doing other things. The organization into processes also allows these to be distributed into different computers to increase capacity and user response times.

The User Database contains a small data base of all the conferences and messages which this user reads. It also contains a small directory of the names (of other users and conferences) most commonly used by this user. In those cases, where the User Database resides on the same computer or LAN as the main data base, storage can be saved by not duplicating all information to each User Database.

3.2. An extensible, object-oriented data base structure

The internal structure of the SuperKOM data base is very simple. It consists of objects which can be of different classes. Each object has the same internal format. Each object has fields for the following information:

- A reference to its creator or owner.
- A main reference to another object, which this object refers to.
- A list of references to objects referring directly or indirectly to this object.
- A creation date and time.
- A list of members of this object.

- A text area, in which message texts and other mostly textual information is stored in an ASN.1 structure.

The most common classes of objects are shown in table 8.

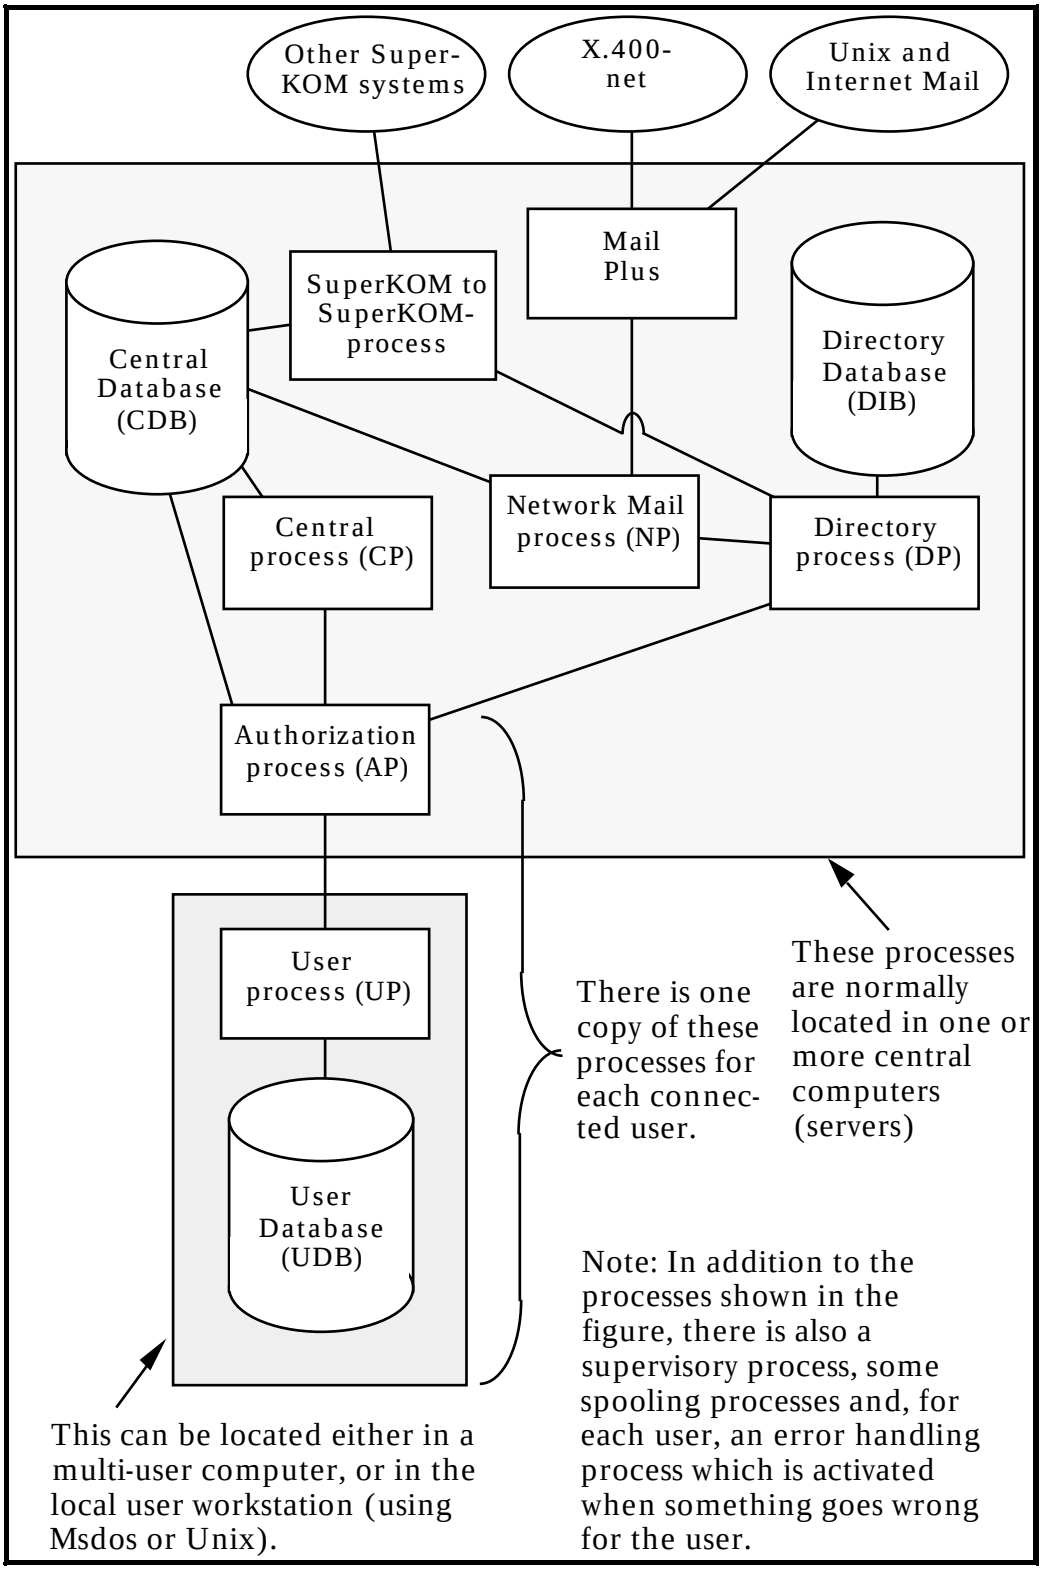

Figure 7
Processes in SuperKOM

Object class	Owner reference	Main reference	Reference list	Member list	Text area
Message	Originator	Recipient or conference	Other messages, which directly or indirectly comment on this message	Normally empty, can contain users who want to read this branch of a conversation only	Text of the message, name of the originator and recipient in text format etc.
Comment, Personal Reply, Referencing message, Single Copy and Obsoleting Message, Additional part	See Message	The message, which this message is a reply to or refers to	See Message	See Message	Can contain the text itself, unless this object only adds a link to another object whose text is stored separately
Conference	Moderator	Moderator	Messages contributed to this conference	Members of this conference	Textual pre-sentation of the conference topic
User	User	User	Personal letters to this user	Empty	Empty

Table 8
Classes of objects in the SuperKOM data base

The data base of SuperKOM can be seen as consisting of objects and links between objects. However, as is seen from the table above, a general class is used which can be both an object and a link at the same time. This is mainly to save disk accesses, since a user normally needs an object and its main links at the same time.

Note the similarity between the structure of a conference and a message. Both can have a reference lists of messages, and a member list of members. This is necessary, since we allow users to subscribe to not only conferences, but also to any conversation or conversational branch. One could say that in SuperKOM, every message is a conference of itself. This does not require much additional disk space, since the member list is empty for most messages, and the reference list is usually very short.

3.3. Information flow when distributing new messages

The distribution of messages in the SuperKOM data base uses a flow model, as shown in figure 9. The act of performing this distribution is done by the background Central Process (CP).

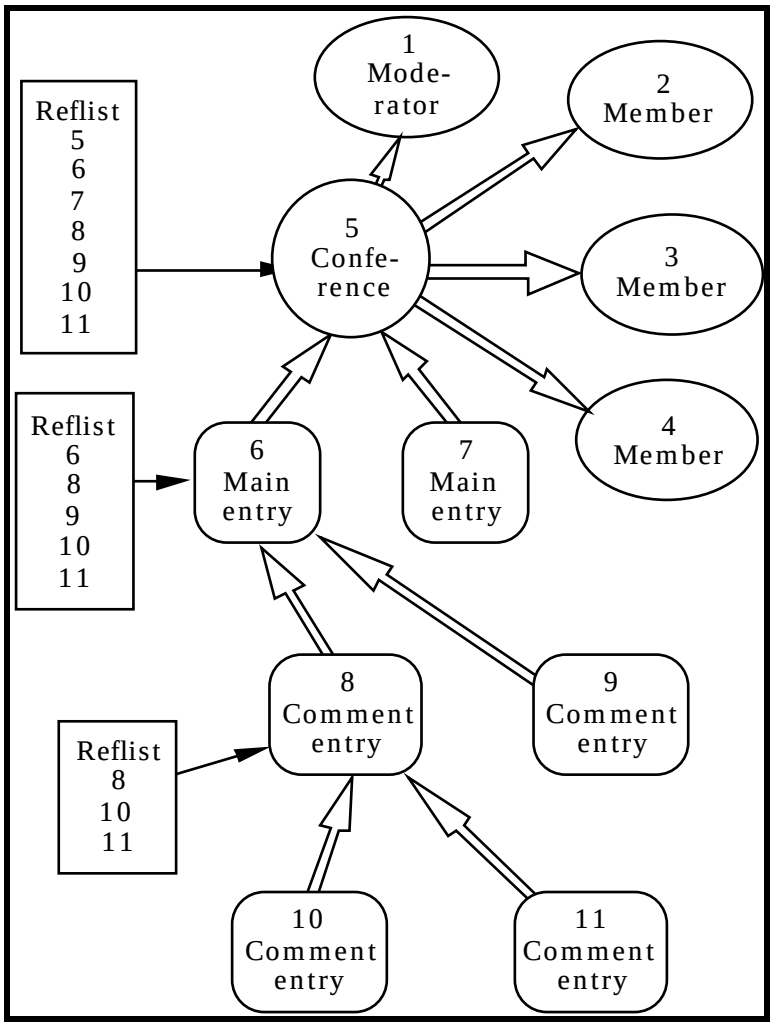

Figure 9
Flow model for news distribution

Figure 9 shows a conference with a moderator, a few members, two main entries, and three entries which are direct or indirect comments on one of the main entries. News distribution is done by letting entries “flow” along the thick arrows representing the links in the data base. Thus, entry 10 flows to entry 8, then flows on to entry 6, then flows on to conference 5, then flows on to users 1, 2,3 and 4, the moderator and the members of conference 5.

Each object keeps a list of the objects which has flowed past it. This list is called a reference list. For a conference, the reference list contains all entries in the conference. For an entry, the reference list contains all entries in this entry and in direct or indirect comments on the entry.

If a new user becomes a member of conference 5, all entries in its reference list will “flow” to the new member. And any forthcoming entry in this conference will “flow” to the conference and then along to the new member.

Since all entries keep a reference list, just as conferences, any user can become a member not only in a conference, but also in any entry. And becoming a member of an

entry means that the user will receive all existing and future entries in the conversational branch headed by that entry.

3.4. Subtractive links

There are two variants of each link, the link itself and its subtractive link. The subtractive link inhibits the link it is subtractive of. Thus, for example, a person withdraws from a conference, not by removing the member link from the person to the conference, but by adding a subtractive link from the person to the conference.

Figures 10 and 11 shows two examples of uses of such subtractive links.

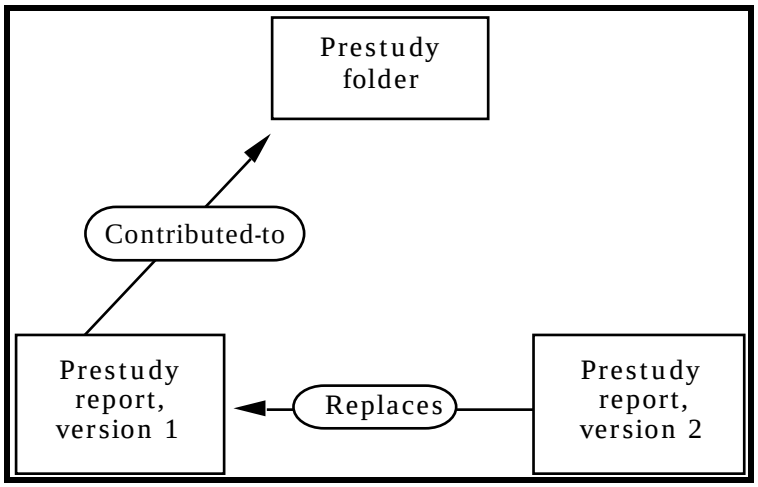

Figure 10
Obsoleting an entry with a replacing entry

Figure 10 shows how an obsoleting link can replace an entry with a new version. Interesting to note is that this functionality, where the old entry can be kept in a data base even after the new entry has arrived, was already hinted at in the 1984 version of X.400 which said (clause 4.2.4.7) “The action to be taken by the recipient or his IPM UA is a local matter. The intent, however, is to allow the IPM UA or the recipient to, for example, remove or file obsolete IP-messages”.

It was probably felt, by the creators of X.400, that a user should have control of his own UA data base, and that the user may not wish other people to be allowed to remove entries from his UA data base by obsoleting them.

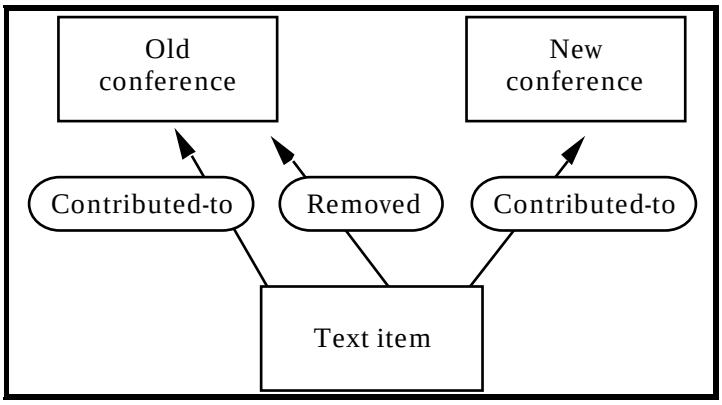

Figure 11
Moving an entry from one conference to another

Figure 11 shows how an entry which is not suitable for one conference, can be moved, by the moderator, to a new conference, by subtracting its link to the old conference with a subtractive link, and instead linking the entry to the new conference.

Note that since comments on an entry are only connected to a conference indirectly via the comment links (as shown in figure 9), moving of one entry to a new conference, as is shown in figure 11, also moves the whole conversational tree started by that entry, including not-yet-written forthcoming entries in that tree. This is useful, to make it easy for the moderator to move a whole branch of a conference to a new conference, or to start a new conference and move a branch from an old conference to the new conference.

There are many advantages with such subtractive links:

- Keeping of the old links even when they have been subtracted keeps a kind of history of previous stages of the data base. It is for example possible too find out which people have been members of a conference, even after they have withdrawn.
- A main idea of computer conferencing is that many users have a common view of the organization of data. Thus, when the data base is changed (for example the moderator moves entries to reorganize the conference structure) this change should be reflected for all members of the conference. But on the other hand, a user should have control of his own mailbox data base, at least to the extent that information is not removed from it without his consent. The handling of deletion of entries not by actually deleting them, but by adding a new obsoleting link to the old version, is an example of this envisaged already in the 1984 version of X.400 (see above). It is up to the user, and to the designer of his UA software, how this is actually shown to users. In SuperKOM, obsoleting entries are handled in the following way:

  (i) If the user has not already seen the old entry, the user will be shown only the new version, but a line in the header will tell the user that it does obsolete an old entry, and the user can find the old entry with a command to the system.

  (ii) If the user has already seen the old entry, the user will be shown also the new version, again with a heading line indicating that it does obsolete an old entry.

- Subtractive links can be used to subtract also indirectly derived links.

  For example (see figure 9) a person who is a member of conference 5, may not wish to read the entries in the subdiscussion headed by entry 8. The user then adds a subtractive link to himself on entry 8. The distribution process will then distribute the entries in the branch headed by entry 8 to all members of the conference except this user. The user command to do this is called “skip all comments” since the effect, as seen for the user, is that the user skips all comments on a certain entry. Note that this command will skip also future, not yet written entries in this conversational branch. There is however also, in SuperKOM, another command to skip only existing entries in a branch without skipping forthcoming entries in that branch.

3.5. A refinement of the flow model for message distribution

The handling of news distribution in SuperKOM is actually a little more complex than described by the flow model in section 3.3 above. Each object is seen as if it has three rooms, the High room (text room), the Middle room (member room) and the Low room (moderator room), as is shown in figure 12.

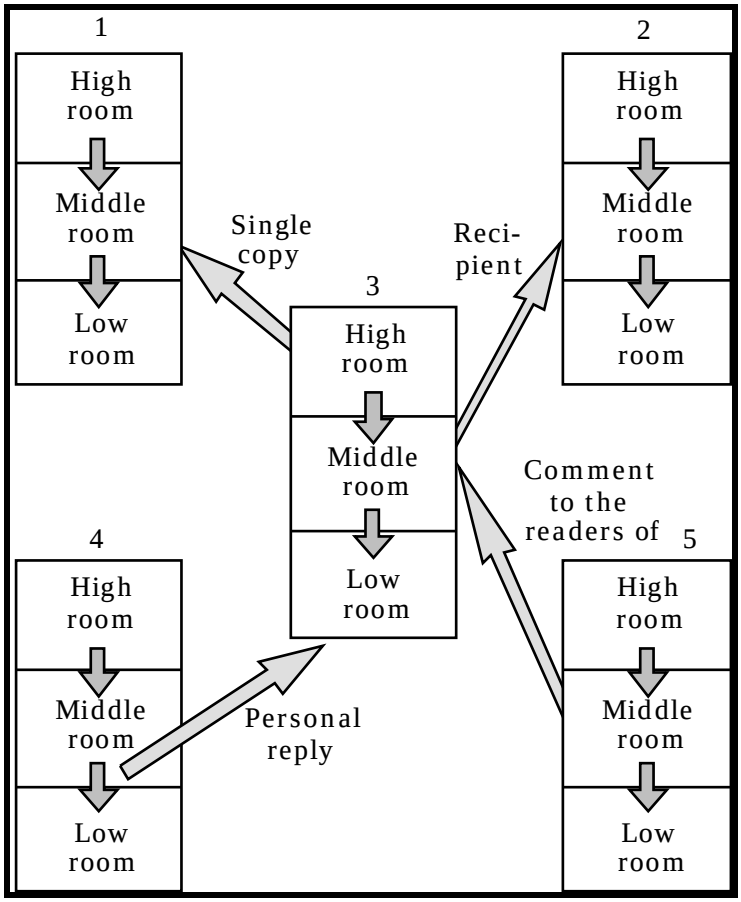

Figure 12
The refined flow model

Each object, as shown in figure 12, is split into three rooms. There are one-way valves between the rooms, so that all objects in the high room will flow to the middle room, and all objects in the middle room will flow to the low room, but not the opposite.

The text of an object itself always starts in the high room and flows on to the two other rooms of that object.

Suppose now that object 5 in figure 12 is a comment on object 3, which is to be seen by all recipients of object 3. The news distribution from object 5 to object 3 then goes between the middle rooms of the two objects, since all normal recipients of object 3 will get the contents of the middle room of object 3.

Suppose that object 4 in figure 12 is a personal reply on object 3, which is only to be seen by the originator of object 3. The personal reply link goes from the middle room to the low room, and only the originator of object 3 will get the contents from its low room.

Object 3 has a recipient link to object 2, which means that all items in the middle room of object 3 will flow on to the middle room of object 2. This means that not only object 3 itself, but also all direct and indirect comments on object 3 will flow to the middle room of object 3 and on to object 2.

Object 1 is a single copy to object 3. By a single copy is meant a link which only sends object 3 itself, but not all comments on it. So the single copy link in figure 12 goes from the high room of object 3 to the middle room of object 1. Since comments on object 3 do not get into its high room, they will not reach object 1.

4. Summary and conclusions

This paper shows that a number of useful functions in a modern conferencing system can be implemented within a rather simple basic structure. Thus, the different features will not be added baubles on a Christmas tree, but rather will naturally derive from the basic structure of the system.

5. Acknowledgements

Ideas have been taken and experience gained from many previous systems, including COM [9], MTS-CONFER [7], COSY [17], EIES [6], EIES 2 [16] [19], MULTICS FORUM [3], PORTACOM [9], PARTICIPATE [5], PLANET-FORUM [18], USENET NEWS [4].

The facilities described in section 2.1, 2.2, 2.3 and 2.9 above were previously available in KOM. The idea for the facilities described in section 2.5 comes from Participate. The idea for the facilities described in section 2.6 comes from COSY and MTS-CONFER. The idea for the facilities described in section 2.7 comes from EIES 2. The idea for the facilities described in section 2.8 comes from Usenet News. The idea for the facilities described in section 2.11 comes from the IBM SAA/CUA specifications, which in its turn has taken the ideas from the Macintosh user interface.

The most important facility of some previous systems which we have not chosen to implement in SuperKOM is a user interface which can be reprogrammed for different applications using a specially designed top-level language. The experience is that in practice such a facility is not very often used.

6. References

[1] CCITT Recommendations X.400-X.430, 1984 and 1988: *Data Communication Networks, Message Handling System*.

[2] Hiltz, S.R. and Turoff, M.: *Structuring Computer-mediated Communication Systems to avoid Information Overload.* Communications of the ACM, July 1985, pp 680-689.

[3] Honeywell: *Multics Forum Interactive meeting system, users' guide.* Honeywell 1982.

[4] Horton, M.: *Standard for Interchange of Usenet Messages.* Network Working Group RFC1036, December 1987.

[5] Keehan, Michael T.: *The Participate computer conferencing system.* AFIPS Office Automation Conference, Los Angeles, February 1984.

[6] Kerr, Elaine B. et al: *Users' manual for the Electronic Information Exchange System.* New Jersey Institute of Technology, 1982.

[7] Michigan University 1985: *Outline of some confer features*. Anonymous.

[8] Palme, J 1984: *You have 134 Unread Mail! Do You Want To Read Them Now?* In Proceedings of the IFIP WG 6.5 Working Conference on Computer-Based Message Services, 1984.

[9] Palme, J.: *Data Base Structure in PortaCOM*. In Byte Magazine, December 1985.

[10] Palme, J 1986: *Cost-Benefit Analysis of Computer-Mediated Message Systems.* In INFORMATION PROCESSING 86, H.-J. Kugler, Ed., Elsevier 1986.

[11] Palme, J 1988: *Two techniques for a distributed message data base.* In Research into Networks and Distributed Applications, EUTECO'88, Rolf Speth, editor, Elsevier April 1988.

[12] ISO 1988: *Group Communication Functionality.* ISO/IEC JTC 1/SC 18/WG 4 N 946, September 1988.

[13] Smith, Hugh, Benford, Steve, Shepherd, Alan and Howidy, Howidy: *OSI-based Group Communication.*

[14] Smith, Hugh, Onions, Julian and Benford, Steve: *Distributed Group Communication: The AMIGO Information Model*, Ellis Horwood, 1989.

[15] Turoff, Murray, Rao, Usha and Hiltz, Starr Roxanne: *Collaborative Hypertext in Computer Mediated Communications.* Proceedings of the Twenty-Fourth Annual Hawaii International Conference on System Sciences 1991.

[16] Turoff, Murray: *Computer-mediated Communication requirements for group support.* Journal of organizational computing, volume 1, number 1, 85-113, 1991.

[17] University of Guelph: *CoSy computer conferencing system, user's guide.* Guelph, Ontario, Canada 1984.

[18] Vallee, Jaques: *The Forum project - network conferencing and its future applications.* Computer Networks, 1(1976) pp 39-52.

[19] J. Whitescarver et al: *EIES 2. A network environment for computer supported collaborative work.* Proceedings of the ACM SIGCOMM Workshop: Frontiers in Computer Communications, August 1987, 230-244.

Jacob Palme is non-tenured professor of computer science at the university of Stockholm and KTH. He has been involved with the development of the COM and SuperKOM computer conferencing systems, and has participated in the ISO and CCITT standards groups responsible for the X.400/MOTIS standards.

Torgny Tholerus is the main designer of the COM and SuperKOM computer conferencing systems. He has been working with COM and SuperKOM development since 1976, and was previously involved in artificial intelligence research at Uppsala University.