The Applications

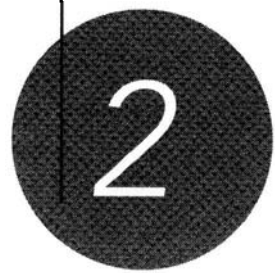

The global Internet liberates many activities from geographical constraints. It joins mass transportation and telecommunications in giving people and organizations considerably more freedom in geographical location while they continue their essential functions. Only a century ago the essence of an organization was centralization—to enhance internal communication—while today, organizations are largely freed from this constraint [Cai97]. An enterprise can be global in extent: A university need not be confined to a campus, a library no longer needs a building, and a community no longer presumes a geographical boundary. All this presumes appropriate applications supporting the necessary activities.

The word "virtual" is often used in conjunction with these applications. A dictionary defines virtual as "being such in essential character." What the Internet—and its close cousins, modern transportation and telecommunications—portends is enterprises, organizations, and communities that are freed of geographical boundaries but are virtual in that they appear to have many desirable properties of geographical centralization.

2.1 Users, Organizations, and Applications

A computing application performs some function on our behalf involving computation, manipulation of information, or communication. It does our work, enlightens us, entertains us, or connects us to other users. A networked application is partitioned across two or

more computers, leveraging the network for communication among application elements.

Two fundamental types of networked applications are those that empower users (people using computers) and those that empower organizations (collections of users with a common mission) and enterprises (organizations with a business mission). There is commonality, in that users may be supporting an organizational mission.

2.1.1 Before Networking

Centralized, time-shared, and decentralized phases of computing have emphasized solitary user-oriented applications. These personal productivity applications enhance the speed or effectiveness of users.

EXAMPLE: The word processor helps users author documents; the spreadsheet automates otherwise tedious computations, allowing many what-if scenarios. Drawing editors turn users into draftspersons (although probably not graphic artists). Speech recognition automates dictation, previously available only to professionals with support staffs.

In organizations, the centralized and decentralized computing eras have been marked by departmental applications. These compartmentalized, hierarchical departments have a specific mission (such as inventory or payroll or purchasing). Prior to networked computing, computing applications were typically retrofitted into the existing departmental structure, serving to extend the capability, improve the productivity, or enhance the quality of its work (much as a personal productivity application extends an individual user).

EXAMPLE: Generating payroll or keeping track of inventory or customers and generating mailing lists are examples of departmental applications.

In contrast to personal productivity applications, the emphasis of departmental applications is often the worker serving the application (rather than the application serving the worker).

EXAMPLE: In inventory tracking, workers may input data, such as informing the application when material arrives to join the inventory.

2.1.2 After Networking

The transition to networked computing shifts the emphasis in user-oriented applications from personal productivity toward social applications, which serve a group of users with a shared mission in some collaborative or communicative activity. In addition, networked computing provides the same capabilities to users who are traveling (called nomadic users). The network supports communication among these users, and networked applications also serve to coordinate their activities. The network also supports virtual interest groups of users sharing common concerns and creates a new public place without geographic boundaries, called cyberspace.

Networked computing is also ideal for creating, accessing, and manipulating information. In information management applications, the network not only enables a single user to globally access vast information resources but also enables information updated in one place to become immediately available everywhere.

Educational applications mix the social aspect (teacher and students) with information access. Business applications serve to coordinate the myriad resources and activities required to produce a product or service and support the communication among workers and managers. Increasingly, business applications extend outside the enterprise, to support and coordinate suppliers and customer relationships—as well as to sell to and support individual consumers—in electronic commerce.

To some extent, both social and business applications predated networked computing. It is feasible to centralize the informational needs of a business process in a single computer, although this has serious disadvantages, including administration (Chapter 3) and scalability (Chapter 10). A time-shared computer can support social applications, as the users sharing that computer can communicate and coordinate. These possibilities are, however, quite limiting.

In contrast, the global Internet supports communication among all users and all organizations. This removes the computing capacity

and administrative boundary limitations of previous phases of computing and enables social applications across organizational, geographical, and political boundaries. The biggest revolution is the virtual interest group or virtual organization. In the virtual interest group, users opportunistically coalesce around a cause or issue. It need not be formally organized or officially recognized (which is why it is called virtual) nor abide by any limitations imposed by geography.

EXAMPLE: Active virtual interest groups have coalesced around issues relevant to the Internet, such as the trade-offs between free expression and limiting children's access to objectionable material, or the trade-off between unfettered electronic commerce and individual privacy. Groups have access to discussion forums and Web pages that link to one another, among other assets that enhance their interest group.

The global network allows any individual or organization to publish information, which is immediately accessible to all citizens with network access. Educational resources can be accessed by all, with fewer geographic limitations.

In business, the Internet allows smaller firms to band together to form virtual enterprises for developing and marketing new products. The virtual enterprise offers many advantages previously afforded large vertically integrated firms, while benefiting from specialized and nimble small firms.

2.2 Application Building Blocks

The building blocks of networked computing applications include physical elements (user, computer, network, etc.) and logical and informational elements (data, infrastructure and application software, etc.). The physical elements are illustrated in Figure 2.1 and briefly described in Table 2.1.

Information is the basic commodity of computing applications. Computers and networks are able to store, manipulate, and communicate digital information (see Section 1.3 on page 9). This information can assume many forms depending on application needs, including

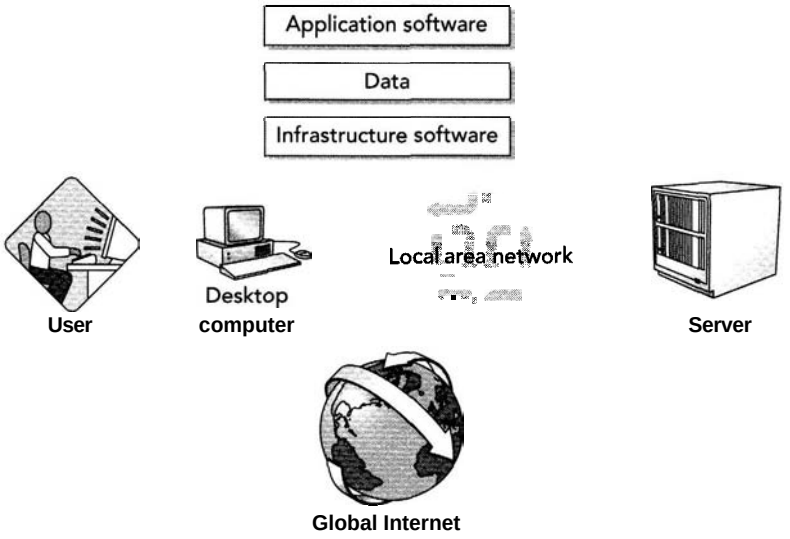

Figure 2.1 The building blocks of networked computing applications.

- Numerical values and text.
- Images: pictures captured by photographic cameras, and their three-dimensional counterpart, virtual reality.
- Graphics: an artificially generated image created by a computer program.
- Audio: sound captured by a microphone.
- Synthesized audio: artificially generated by a computer program.
- Video: a sequence of images representing time as well as space.
- Animation: video artificially created by a computer program.

Applications also manipulate information in ways that combine these media. A premier example is the document—often used to archive or communicate human knowledge—which can combine text, images, and graphics in various ways. Computer-mediated documents can be multimedia, meaning they combine the usual features of paper documents with audio, video, animations, etc. All these forms of information, including audio and video and multimedia documents, can be communicated over the network.

Table 2.1 Description of the building blocks shown in Figure 2.1.

Building block	Description	Typical function
User	Person who interacts with and derives benefit from a networked computing application.	Enters, retrieves, or manipulates information, or performs calculations.
Desktop computer (or personal computer or client computer)	A computer directly accessed by the user that acts as an interface between user and application.	Its screen displays windows, menus, graphics, graphs, etc., and its keyboard, pointing device, camera, and microphone accept input from the user (Chapter 3).
Server computer	A computer not directly associated with a user, usually missing a display, keyboard, or pointing device (except perhaps for administrative purposes). Often has substantial computing power and storage capacity and peripherals.	Stores, accesses, and manipulates large repositories of data, and realizes the logic of how the data is manipulated in light of user directives (Chapter 3).
Application software	The programs running on the clients and servers realizing application functionality.	Embodies the unique functionality of the application (Chapters 4 and 6).
Data	The collection of bits representing—within the computer and network—the information manipulated by the application.	The form in which information is stored, processed, and communicated (Chapter 6).
Infrastructure software	The programs running on the computers and providing many common needs for all applications.	Provides communications support (Chapter 7) and manages resources (e.g., storage and memory-see Chapter 10).
Network	The communication infrastructure connecting the computers running an application. This may be a local area network (LAN) or a wide area network (WAN) such as the global Internet.	Allows the programs running on the different computers to communicate data (Chapter 1.1).

The major categories of applications supported by networked computing can be divided into general classes, including social, information management, educational, and business applications.

2.3 Social Applications

Social applications focus on supporting activities of groups of users, whether or not a group is associated with an organization or enterprise. They are sometimes called "collaborative applications." although collaboration is only one of numerous activities supported. You are no doubt familiar with many of these applications, including telephony (two users holding a conversation), email and voicemail, newsgroups, and chatrooms. The possibilities are much richer than suggested by these early successful examples. Social applications can be categorized not only by the characteristics of the group of users participating but also by what the group is attempting to accomplish.

2.3.1 Characteristics of User Groups

Before considering many representative networked applications, it is helpful to classify groups of users according to their characteristics. The most pertinent group attributes include

- Number of users: A group served by a social application can range from two users to the entire population of users with network access.
- Narrowness of purpose: Some groups form for the purpose of accomplishing a specific task (like scheduling a meeting), while other groups are very unfocused (like those coalescing around a discussion on some topic).
- Duration: Some groups only exist for a short period of time (like two users participating in a telephone call), while others can be very long lasting (like the coauthors of a book).
- Social relationships: In some groups—particularly small ones—all the users know who the other users are and may well know them personally. In very large groups, each user will typically not even be aware of who the other users are, let alone know them. In the

latter, there may be subgroups who are mutually familiar or friendly, or social relationships may be more diffuse.

With these group characteristics in mind, a classification of the types of groups supported by social applications is shown in Table 2.2. The terminology defined in this table will be used in the remainder of this chapter.

As implied in Table 2.2, these types of groups actually have a structural relationship like that of Figure 2.2. The citizenry includes a very large group of users, subsets of which form ongoing interest groups. The number and composition of these interest groups

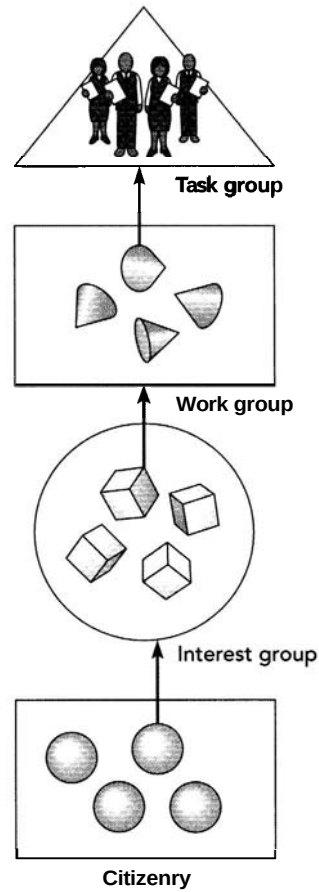

Figure 2.2 The inclusive relationship of group categories.

Table 2.2 A classification of groups using a social application.

Category of user group	Description	Examples
Individual	A solitary user working to accomplish some goal. This is the limiting case of no group. Computer applications are meant to enhance quality and productivity.	A user writes a single-author memorandum; a user adds an appointment to his personal calendar; a user accesses stock prices using the Web.
Task group	A task is a short-term effort directed at an immediate goal. Members of a task group interact with one another to complete a task. This interaction may require the undivided attention of all users.	One user telephones another to make a lunch date; a group of workers holds a meeting to evaluate a competitive bid or plan the next step in a project.
Work group	A project is a longer-term effort directed at a challenging goal. Members of a work group collaborate to complete the project. A project may spin off short-duration tasks addressed by constituent task groups.	Employees from the marketing, sales, engineering, and finance departments develop a product plan; scientists write a joint-authored paper; members and staff of a legislative committee work on a new computer security bill.
Interest group	An interest is a topic, profession, hobby, or goal that is ongoing and open-ended. An interest group (occasionally called a community) pursues an interest through common discussion, study, or collective action. Interest groups typically form constituent work groups to collaborate on projects related to the group interest.	Historians share a common interest in World War II; world's coin collectors pursue an interest in coin sale and valuation; employees of General Motors Corporation pursue an interest in building and marketing automobiles; users concerned about network privacy pursue an interest in policies and laws.
Citizenry	A large group of users without a specific organized purpose. Typically subsets of the citizenry form ongoing interest groups around topical issues.	Citizens of San Francisco, California; all the users with Internet access or telephones.

change with time, but on a timescale of weeks or months. Similarly, interest groups break off to form sporadic work groups to pursue particular projects (such as the design of a product or the organization of an international conference). In turn, a work group may require sporadic task groups that interact on short-term tasks, interspersed with the individual effort of members of the task group.

EXAMPLE: American business is a citizenry, and an interest group within is the employees of General Motors (GM). When GM designs a new car, it forms a team (work group) that collaborates on the design project. That collaboration consists, in part, of numerous short-term tasks that are addressed through interaction in a meeting or conference call. A typical project is to develop a marketing plan—involving marketing, sales, development, manufacturing and distribution, etc.-with a detailed planning document as the outcome. This may take several months, during which there could be task groups formed to interact on particular issues (such as developing a schedule and coordination plan for the next month).

Table 2.3 lists typical characteristics of these categories of groups (for the three largest groups). Although each user likely participates in at most one task group at a time and gives it her undivided attention, she may participate in multiple work groups and interest groups and split her time and attention among them.

2.3.2 Styles of Social Applications

Having categorized the groups participating in social applications, the next step is to appreciate the characteristics of the applications themselves. Examining current and emerging applications, there are two major styles (which will subsequently be broken down further):

- Communication: Any group effort requires the communication of information among group members. This may take the form of a discussion or the circulation of draft documents.
- Coordination: To accomplish its goals efficiently and effectively, the different activities of a group must be coordinated. Often, certain tasks or projects depend on the outcome of other tasks

Table 2.3 Typical characteristics for different categories of groups.

Group category	Work group	Interest group	Citizenry
Number of users	Small, typically from two to ten or twenty.	Large, typically from hundreds to tens of thousands.	Everybody.
Narrowness of purpose	Typically pursues a specific near-term project goal with a defined outcome.	Although constituent interactions and collaborations are focused around the mutual interest, they may take the group in many different directions (even simultaneously).	No specific common purpose or interest.
Duration	Typically days to months.	Indeterminate, as there is no predefined goal.	Forever.
Social relationships	Each user typically knows who the other group members are and may personally know many of them.	Each user typically knows only a small subset of the other group members.	Each user knows a tiny fraction of the other group members.

or projects or individual efforts, so the proper ordering must be coordinated. The group may require common resources requiring coordination, such as preventing conflicts in the joint editing of a document. In addition, communication opportunities must also be coordinated, for example, scheduling the time of a meeting.

'Together, communication and coordination are what distinguish social and isolated activities, and the objective of social applications is to improve and enhance both dimensions — even for a group that is geographically or administratively dispersed. Both styles are now examined further.

Communication Style

The communication style of social applications can be further divided into four substyles, as shown (with example applications) in Table 2.4. One style dimension accommodates the degree of knowledge that one user has of other users:

Table 2.4 Examples of communication applications in each of four substyles. Applications in each row are immediate and deferred variations.

Social application styles	Immediate (users participate simultaneously)	Deferred (users need not participate simultaneously)
Direct (users know precisely which other users are participating)	Telephony and video conferencing simulate "face-to-face" interactions of users (a meeting).	Electronic mail (email), voicemail, and facsimile allow one user to originate a communication and another user to access it later.
Publication (users do not know other participating users)	Broadcast video (analogous to broadcast television) allows one user to simultaneously address many other users for seminar viewing or distance learning.	Video on demand (analogous to the video rental store) allows one user to store a video presentation on a server and other users to view that video at a time of their choosing.
	Information push (analogous to a newspaper) allows one user to publish volatile information (like stock quotes) to be viewed immediately by other users.	World Wide Web (an example of a more general category information pull) allows any user to publish information to be viewed later by other users.

- Direct style: In the task group and work group, users typically know the other users in the group. This admits a direct style in which users communicate directly with other users.
- Publication style: In the interest group and citizenry, direct communication may not be possible—because the users don't know one another—except of course when the group forms smaller work groups or task groups. In spite of these loose or nonexistent social relationships, communication within the group is valuable, for example, to disseminate ideas or form work groups. In the publication style, one user (or small group) makes information available in a form that can be accessed by any other user. By its nature, a publication benefits the group as a whole, not

specific users. Each user makes his or her own decision to access the information, and typically some do and many do not. The user publishing the information cannot anticipate who will eventually access it.

ANALOGY: The direct style is analogous to how the telephone network is used, or a memorandum in business with a specific distribution list (an important tool for work groups). Scholarly journals (an important tool for scholarly interest groups) and newspapers (an important tool for the citizenry) illustrate the publication style.

The direct style of communication distributes information to a known set of recipients (possibly a single recipient), while the publication style distributes it to an unknown set of recipients. The direct style is most appropriate for task groups and work groups, while publication serves interest groups and the citizenry.

The other style dimension for communication applications makes a distinction based on whether users participate simultaneously or not:

- Immediate style (sometimes called synchronous): In this style, users participate in the application at the same time. This is practical for a task group, where the number of users is small and they know one another. It is not practical for the citizenry, because it would be impossible to schedule them to interact at the same time, and there would be too many participants for this to be effective.

  Deferred style (sometimes called asynchronous): This style removes the constraint that users participate simultaneously. This eases scheduling difficulties, reduces the invasiveness of the communication to the individual user, and increases the size of the group for which communication is feasible.

ANALOGY: A mother phoning her daughter to wish her happy birthday is the immediate style, while sending a birthday card in the mail is deferred.

Coordination Style

Group activities create dependencies among users and other resources. For example, the very viability of immediate applications makes the users dependent: They have to participate at the same time. Often, deferred communication applications create similar dependencies. For example, users collaboratively editing a document have to work on it in a particular order, they have to avoid making conflicting changes, etc. Coordination style applications manage these dependencies, expediting the completion of a task or project. This style of application can be further broken down into two substyles, as listed in Table 2.5.

Coordination applications particularly aid work group project management and allow workers to minimize disruptions caused by conflicts or delays in receiving necessary resources.

How do these application styles relate specifically to the needs of groups? The following subsections (and sidebars) describe a num-

Table 2.5 Two styles of coordination application.

Style	Description	Examples
Resource allocation and scheduling	The members of a group share resources, which must be managed for efficiency and to avoid conflicts. One aspect of resource allocation is the scheduling of a shared resource so that it can be used by different users or groups at different times.	A meeting room must be scheduled so that only one meeting occurs at a time. In the collaborative authoring of a document, the additions or changes of different users must not conflict. Members of a task group participating in an immediate application must be scheduled. Auctions or other economic mechanisms can be used to allocate consumable resources (see the discussion of electronic commerce later in the chapter).
Monitoring and notification	Monitoring (sometimes called awareness) applications allow group members to benefit from information about some remote resource or user. Notification provides an alert that some condition has occurred.	One user can monitor the availability of another worker for a direct/immediate interaction (like a telephone call). A work group member can request notification when a conference room becomes free. The productivity of workers can be monitored (this raises privacy concerns—see [Gar89]).

ber of social applications, in each case relating them to the group categories and application styles. Collectively, these applications illustrate the wealth of valuable networked applications. They are meant to give you the sense that the space of possible applications is very rich, with many unexplored possibilities.

2.3.3 Remote Conferencing with Shared Workspace

Without networked computing, a task group or work group might find a conference room to conduct their collaboration. In that conference room, they would hold a meeting—with face-to-face discussions—and also locate any work items to be examined and modified (such as documents being collaboratively edited). The team would likely also use a visual aid such as a whiteboard to share ideas.

Such an interaction or collaboration can also be conducted over a large geographical area by using a social application called remote conferencing with shared workspace. This direct application serves task groups or work groups. The direct-immediate form of this application attempts to reproduce all the facets of a physical conference room, including

- Telephony: Speech is the most basic form of human communication, and thus telephony (holding a conversation at a distance using speech) is the most successful communications application. Telephony is provided by the telephone network but can be provided in a networked computing infrastructure as well.
- Video: Humans also communicate through facial expressions and gestures, and thus a video presentation of remote users can lend a feeling of presence, proximity, and trust that contributes to the quality of the interaction. Video conferencing is a combination of telephony and video and can be enhanced by other media, such as those listed next. (However, many feel that video is the least important element of a conference.)

- Presentation graphics: It is common in meetings to use visual aids like slides and transparencies. Since they are prepared in electronic form, they can be projected to remote users.
- Collaborative authoring of a shared document: A document that is being collaboratively authored can be stored somewhere and the group members allowed to view and edit it (see the sidebar "Collaborative Authoring").
- Hand drawings and doodles: Participants in a meeting frequently communicate ideas or designs through hand drawings on a whiteboard or blackboard. These drawings can be captured and communicated to remote users, as they are drawn, using a liveboard. Alternatively, a shared whiteboard application uses a mouse or tablet to draw, with the result displayed remotely.

EXAMPLE: The screen capture shown in Figure 2.3 shows the visual aspect of a remote conference. Shown are standard applications developed for the multicast backbone (Mbone), which is an Internet capability for sending audio and video media from one source to many destinations (see Chapter 7).

2.3.4 Groupware

Remote conferencing is primarily intended as a direct-immediate application serving a task group, where users work directly together, at the same time, and the collaboration may have their undivided attention. Unfortunately, this style becomes ineffective when a group becomes too large. For example, this makes it difficult to capture and display video of everyone, along with whiteboards, on a small screen. Too many participants also reduce the quality of the interaction (as is starting to happen in Figure 2.3 as the screen gets cluttered).

ANALOGY: Consider the quality of the interaction when a dozen, a hundred, and a thousand people meet face-to-face. A dozen people can have an effective meeting, allowing everyone to participate and interact. On the other hand, a hundred or a thousand people cannot effectively hold an interactive meeting. It takes too long to hear from everyone on any particular point—the number of inputs is beyond the point of diminishing returns—and it becomes difficult

Collaborative Authoring

The most basic function of a collaborative authoring application is allowing any user to view and edit a group document. Because of possible conflicts when two or more users edit the same document, the application also coordinates the users in several ways, for example:

- Access control and locking limit who can edit and who can read documents. Access can be restricted to particular users or temporarily precluded while a document is edited.
- Version control keeps track of current and past versions of the document. For example, anyone can see who made what changes, see what those changes were, undo them, etc.
- Annotation of the document allows one user to pass comments to another user (without editing the document itself). Comments can be attached to the precise location where they apply and can be multimedia (for example, voice rather than text).
- Replication and reconciliation are sophisticated capabilities. Normally, editing a given document would be restricted to one user at a time, slowing the authoring process. Replication creates two or more

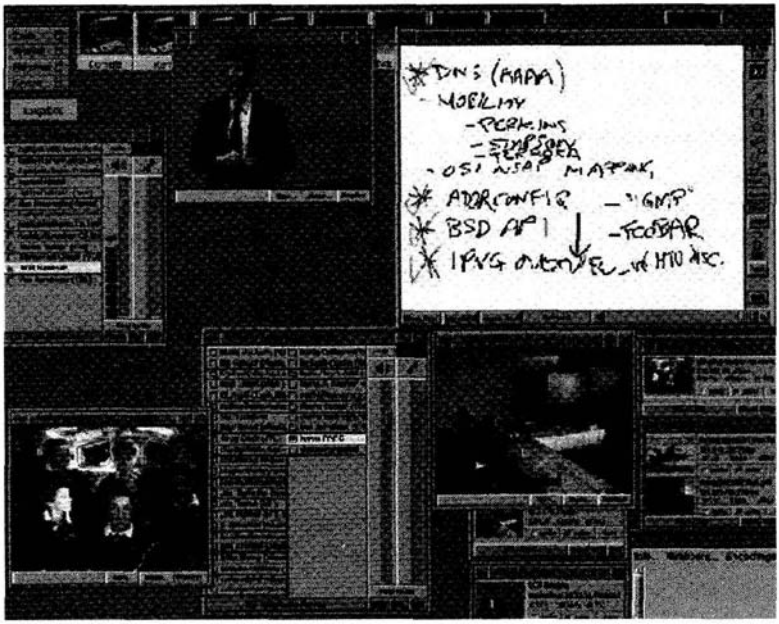

Figure 2.3 Example of several collaborative tools for the Mbone: vic (video), vat (audio), and wb (whiteboard) from [McC95]. Source: this image depicts experimental collaborative applications developed by the MASH Research Group at the University of California, Berkeley. See http://www-mash.cs.berkeley.edu/mash for more information on the MASH Project.

to fairly allocate talking time without hearing the same points over and over.

These problems can be avoided by a deferred version of the remote conferencing application, called groupware. Because it is deferred—allowing group members to participate when they choose—and also because it can aid larger groups, groupware is appropriate for work groups.

EXAMPLE: The major suppliers of groupware are Lotus (a division of IBM), Microsoft, and Novell, and each sells complementary client and server software (Lotus Notes and Domino, Microsoft Outlook and Exchange, Novell WebAccess and GroupWise). These products began with deferred work group and document management

identical replicas of the document that can be independently edited. The problem, of course, is merging the changes back into a single version, which is reconciliation. Reconciliation is relatively simple if changes are made in independent places and more complicated if common sections are modified.

As you can see, collaborative authoring applications mix communication and coordination. The users communicate through the document itself as well as through annotations, but the application also coordinates the shared access to the document.

but increasingly incorporate immediate conferencing capabilities as well. An example is real-time messaging, in which a message appears immediately on group members' screens.

Collaborative authoring — discussed in the previous section and the sidebar "Collaborative Authoring" — is a pillar of groupware. The joint editing and coordination functions associated with document authoring don't depend on the users participating simultaneously-they can easily be deferred as well as immediate.

Groupware has to reproduce — in a deferred style — the communication facilities afforded by remote conferencing (see the previous section). It also has added coordination challenges, since the users themselves perform many needed coordination functions when they are interacting directly and immediately.

The most basic capability in groupware is messaging. Users can send one another a message--a package of information that one user wants to pass to one or more other users.

ANALOGY: In the physical world, when one person is unable to telephone another, she might use a postal letter, which is a form of message. The letter encloses in an envelope whatever information the sender wishes to convey to the recipient. The letter could include not only text but also pictures or audio or video (stored on magnetic tape).

Like the postal letter, a groupware message can include multimedia information, such as text, an audio recording for playback, pictures, graphics, etc. A message can even include an entire formatted document. The recipient can read the text, play back the audio or video, look at the pictures, and read the document. A message is a form of direct-deferred communication between users: The recipient need not participate at the same time as the sender but can look at the message at the time of his choosing. The message replaces the direct interaction via audio, video, images, etc., in remote conferencing.

EXAMPLE: The most basic messaging service is electronic mail (or email). A sender of an email message can designate one or more recipients for that message. Each user has an inbox where mes-

sages arrive and await access by that user. Although only text messages were supported originally, increasingly, email applications allow multimedia messages. Email is an almost direct analogy to the postal service in the physical world, including notification of a nonexistent recipient or notification of delivery and access by the recipient.

As illustrated by email, a message can be designated for a single recipient or multiple designated recipients. Since a given user can receive messages from many senders, a message system must merge these messages. Message applications can add a number of other features, such as

- Priority: Senders can attach priorities to messages, so that recipients can access high-priority messages first.
- Filtering: A fundamental issue with message systems is that a recipient has no control over who can send him a message. This may result in wasted time sifting through many uninteresting messages. A message filter can discard messages that don't meet criteria specified by the user (negative filtering) or, alternatively, only allow messages that do meet specified criteria (positive filtering). These criteria might include the sender's identity or be related to the subject or content of the message.
- Authentication: It might be possible for an imposter to send a message; authentication verifies the identity of the sender.
- Integrity: A message might be modified somewhere between the time it is sent and it is read, either accidently or for some nefarious purpose. A message with integrity is assured to be exactly as composed by the sender.
- Confidentiality: Some messages contain sensitive information that should be available only to the recipient. Confidentiality ensures that only the recipient can read it.

The first two features add coordination functions, and the latter three enhance the utility of the communication (and can just as easily be incorporated in remote conferencing as well).

Aside from collaborative authoring and messaging, groupware also includes coordination capabilities. Even in a work group, messaging

Calendar and Scheduling

Calendar and scheduling applications manage the personal calendar of each user and also schedule users for task groups (meetings, remote conferencing, telephone calls, etc.). They can schedule auxiliary resources (such as meeting rooms or video conference facilities). A user must publish his or her personal calendar for the application—minimally, times available for task group interactions, or more information if automatic rescheduling is desired. The more willing the user is to relinquish personal control, the more automated and effective scheduling can be.

Calendar and scheduling illustrates the tension between automation and privacy. Without the personal involvement of the user, it may become too easy to schedule interactions of marginal value. For the future, ways are needed for users to describe priorities and automate interaction possibilities. Otherwise, users may be reluctant to cede total control to a faceless application.

is rarely sufficient for communications. Sporadic remote conferences and face-to-face meetings are valued, especially for establishing mutual trust and dealing with complex issues or negotiations. Since remote conferencing is an immediate application, it has to be scheduled at a time mutually suitable to members of the task group, and they may even have to rearrange schedules to make an immediate interaction feasible.

EXAMPLE: You have encountered the endless round of voicemail messages often required to talk to someone (called telephone tag) or the time consumed in coordinating and juggling everybody's schedule for a meeting (accentuatedin the global economy, where time zones reduce the feasible times). These problems reduce the viability of direct-immediate applications, in spite of their compelling advantages for task groups.

Calendar and scheduling is a groupware application that eases these logistics by making the calendars of individual users available to a scheduling application, which can access and manipulate them to coordinate schedules (see the sidebar "Calendar and Scheduling"). It illustrates a publication-deferred application, since each user in the group publishes his or her calendar for the benefit of anybody in the group.

2.3.5 Discussion Forums

Typically task group interactions are scheduled with forethought to a specific purpose and agenda. This is fine for well-defined outcomes with clear steps to get there. On the other hand, discussion and creative brainstorming have no predefined outcome or stopping point. They can be performed by a task group but often are more effective in a deferred style. Performing brainstorming over a longer period of time—interspersed with other activities—is often less intimidating to participants and more conducive to new ideas. Brainstorming is supported by the discussion forum, where any group member can propose ideas or comment on ideas previously proposed.

The remote conference is a direct-immediate style of discussion forum that serves task groups. The discussion forum is an even

more important tool for interest groups — for which direct applications like the remote conference are not feasible because users often don't know other users in the group — because it is typically realized in a publication style. That is, a discussion forum works by one user publishing an idea in the forum (by sending a message to the forum application, rather than directly to other users) where it becomes available for any user in the group to access or not access at his or her option. Thus, the discussion forum illustrates that a messaging application need not be direct: A message can be sent to an unknown set of recipients. It is also possible for discussion forums to be anonymous — the sender's identity is not revealed. There are typically many simultaneous discussion forums on different topics, serving different interest groups.

Like most social applications, the discussion forum comes in deferred and immediate styles. A deferred style of discussion forum is the newsgroup (see the sidebar "Newsgroups"). A discussion is started by one user sending a message (called a posting) to a common repository of messages (called a newsgroup) that is published for the benefit of the group.

ANALOGY: As China loosened political control, there emerged a "democracy wall" devoted to discussion of political freedom. This was a physical wall in Beijing, where anybody could post thoughts on paper, and anybody else could read them and post responses.

Chatroorn

Since a newsgroup is deferred, the immediacy of a face-to-face interaction is lost. A remote conference — being a direct style of application — is not a suitable replacement for a newsgroup. This motivates the chatroom — a publication-immediate style of discussion forum.

ANALOGY: The chatroom is analogous to a continuously running town meeting. Any member of the interest group can join the discussion in a town meeting at any time. Unlike a remote conference, the other participants are not known in advance.

Newsgroups

A publication-deferred application, newsgroups associate a topic of discussion with a specific subject heading and are particularly targeted at interest groups. Any user can post a message relevant to that topic, and any other interested user can read previous postings and post responses. Typically, each message posted on a given subject stimulates responses from other group members. Those responses are posted under the same subject heading, and that group of messages forms a thread. The users joining a thread are a working group that collaborate on that specific subject. Note that the thread forms opportunistically about that subject—the user initiating the thread does not have to anticipate which other users may be interested in participating.

The newsgroup has mechanisms for hiding unwanted messages. For example, previously read messages can be hidden, and threads are collapsed under a single heading, with the individual messages visible only if desired. A newsgroup may also have a moderator—a user with special authority to determine which postings are allowed, to change the sub-

Any user can join an ongoing discussion by reading and posting messages to other users who happen to be simultaneously participating. It works like a remote conference associated with a named topic. Users participate by "entering" the chatroom, after which they see all postings immediately as they occur. Anybody with something to say on the topic can post a message, which other participants see immediately. A posting may engender an immediate response from others—hence the spontaneity and immediacy-leading to a "conversation" that can be viewed by all.

EXAMPLE: The group of customers (and potential customers) of a company is an interest group. Chatrooms are used for customer service, allowing customers to communicate with service agents and one another. Acuity Corporation and Business Solutions Inc. are two providers of chatroom applications.

The chatroom can be viewed as either an immediate variation on a newsgroup or as a publication variation on a remote conference. Like other messaging applications, messages in both newsgroups and chatrooms can be multimedia (incorporating audio, video, whiteboard, etc.).

Listserver

With a newsgroup, each user must make a conscious effort to periodically check postings, and a chatroom requires a user's undivided attention to derive full benefit. This does not work well for extremely busy people, or interest groups with infrequent postings. A listserver is a publication-deferred variation on the discussion forum that eliminates the published repository of messages and follows a subscription model, in which the user doesn't ask for specific information, but rather all available information on a specific subject or topic. Subscription is an important mechanism for information access, discussed in Section 2.4 on page 38.

ANALOGY: A special-interest magazine serves an interest group. Each group member subscribes to the magazine, and each issue thereafter appears in that member's mailbox.

A user wishing to join a topic subscribes by providing his or her email address to the listserver application. Any user can post a message to that topic, and each posting is automatically emailed to all subscribers, who therefore do not have to consciously access the postings. A user can also cancel the subscription and leave the interest group.

The email messaging system incorporated into the listserver application merges the messages coming from other users and other listservers. A disadvantage is that a subscriber's mailbox may be inundated with messages—especially as the interest group grows large. Because some users might want to "listen in" without posting, or avoid a large number of messages, a listserver can also maintain an archive of past postings.

ject heading of a thread to make it more transparent, etc. The moderator's job is to ensure order and organization (rather than anarchy and chaos) and keep the discussion on track. A newsgroup may also have searching capability, allowing it to look for threads whose subjects contain specific keywords.

2.3.6 Cyberspace Applications

Societies have always had public places, like the town square or public park. The Internet has created a new virtual public place, popularly called cyberspace. Citizens can go there to interact with others, share ideas and criticize government, or just hang out. Like a public park, crimes can be committed there, or privacy violated, or misleading or inflammatory information distributed. Cyberspace is global and hence not subject to the ordinary geographically based jurisdiction of governments.

Cyberspace citizens form short-term task groups (for example, to find a suitable date for dinner), form work groups to collaborate on a project (for example, to organize a neighborhood crime watch committee), and form interest groups (for example, to run a Boy Scout troop or rally around some cause). Thus, the social applications described earlier apply to the citizenry as well. Like an interest group—but even more so—communication frequently occurs among citizens who may not know one another in advance.

What applications specifically support the citizenry? Thus far, they mostly mirror the physical world.

Broadcasting

As mentioned in Section 2.2 on page 16, the information communicated through the network can be multimedia: It can include and mix text, documents, audio, video, pictures, etc. Cyberspace broadcasting—similar to radio and television broadcasting—similarly sends audio and video through the network to any citizen wishing to listen or watch. The radio spectrum limits the number of different radio and television stations; for example, broadcasters using the radio spectrum must be coordinated (by government licensing) to avoid mutual interference, and the number of licensees is severely limited. On the Internet, no government license is required (at least in most countries), and the restrictions are few (although there are issues such as protecting children—see Chapter 5). One can imagine, for example, all the broadcasts occurring anywhere in the world to be available to any citizen in cyberspace. As the Internet advances, it will accommodate many more broadcasters and many more specialized options.

Mass Publication

In the physical world, citizens are informed about current events by the mass media, including newspapers and magazines. The publication-deferred application supporting mass publication in cyberspace is the Web (see the sidebar "World Wide Web"). Many of the same publications available in the physical world are also published on the Web. Web publication can be multimedia in applications called audio-on-demand and video-on-demand, which are analogous to the video rental store. Broadcasting is a publication-immediate application, and mass publication is a deferred variation. There are variations that mix these styles—called information push—discussed in Section 2.4 on page 38.

In cyberspace, publication and broadcasting are inexpensive, making it possible for ordinary citizens to take advantage of them. All that is required to publish a Web page worldwide is Internet access and a desktop computer. Although cyberspace is increasingly populated by large corporations and commercial activity, this in no way precludes any individual from using the medium to express her or his views. Also, while conventional publishing and broadcasting is predominantly a one-way medium, cyberspace is more democrati-

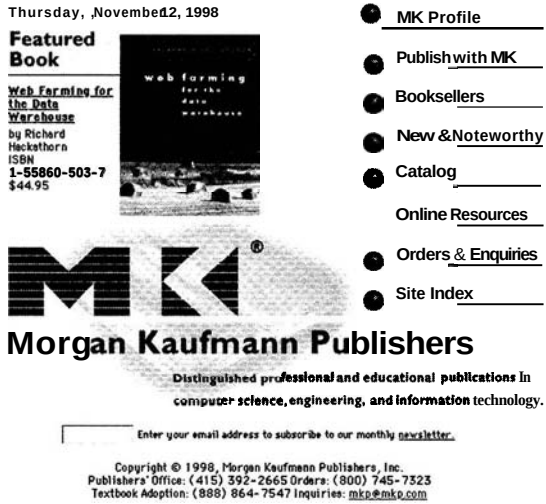

Figure 2.4 A window from a Web browser, showing a typical page with hyperlinks.

cally two-way. It allows various forms of interactive publishing and broadcasting.

Information Retrieval

An overriding issue of cyberspace is, since citizens don't know one another and the volume of information available is huge, how do citizens find useful information or narrow down information to that which is useful? Similarly, how do citizens find and join interest groups?

There are a number of possibilities. The Web supports a hypertext model, which allows users to interactively navigate through large volumes of information by following links from one document to another. There are also search engines that look for information on particular topics. As in libraries and bookstores in the physical world, information indexes are available. In the commercial realm, advertising is an important mechanism for alerting consumers. These and many other possibilities are described in Section 2.4.

World Wide Web

The Web is a publication-deferred application (that has become almost synonymous with the Internet) offering rich possibilities for the publication of multimedia information. A user accesses the Web using a Web browser and sees various text, image, audio, and video pages (see a typical page in Figure 2.4). The browser can pull information from many Web servers holding a variety of information. Initiated by a user, the browser requests a page of information from the server, which is displayed by the browser. Each page can include hyperlinks, which allow immediate access to a related page (by merely clicking on the highlighted hyperlink).

The Web allows any user to publish information in cyberspace, and any other user to access it. The Web also serves an archival function, since published information remains accessible (as long as the publisher chooses).

Consumer Electronic Commerce

In the physical world, citizens go to a merchant's store to buy goods. Similarly, goods can be purchased in cyberspace from merchants who set up "storefronts" there. The Web has evolved into an application that supports consumer electronic commerce as well as information retrieval. If information goods are purchased, they can be delivered in cyberspace as well; otherwise, they may be delivered by the transportation system. Electronic commerce is discussed in Section 2.6 on page 52.

Recommendation Sharing

One way of identifying useful information, products, or interest groups is recommendation sharing. A recommender system is a communication publication-deferred style of application that collects recommendations from many users and makes them available to any interested user (see Section 2.4).

2.3.7 Back to the Big Picture

The specific social applications are listed in Table 2.6 by category of group they serve. Also listed are some analogous mechanisms in the physical world.

2.4 Information Management

One major category of networked application is the storage, manipulation, and retrieval of information. In the physical world, this is the domain of libraries and bookstores, which house vast information resources and offer ways to identify useful or targeted information. However, as the volume of information grows, libraries exhaust their physical space. The physical shelves of the library can be replaced by digital information repositories (residing in computer storage), in which information is represented by data (a collection of bits—see the sidebar "Any Information Can Be Represented by Bits" on page 10). The manipulation and searching of information can be performed by networked computers that have access to those repositories over the network, often guided interactively by users or librarians.

Table 2.6 Social applications organized by target group category.

Group category	Social applications	Physical-world analogy
Individual	Word processing; spreadsheet	Pencil and paper; electronic calculator
Task group	Email and voicemail; telephony and video conferencing; whiteboard	Postal letters; telephone; facsimile
Work group	Remote conferencing with shared workspace; calendar and scheduling; collaborative authoring; monitoring and notification	Face-to-face meeting; administrative assistant
Interest group	Discussion forum: newsgroup, listserver, chatroom	Town meeting; special-interest magazine
Citizenry	Publication: broadcast, Web; consumer electronic commerce; recommendation sharing	Television and radio; newspaper; news magazine; catalog shopping

Compared to the physical alternative, networked computing has some compelling advantages in its ability to store and manage information:

- It is especially effective in accessing volatile information—that which changes frequently (e.g., stock prices). Those changes-entered in one place—can be reflected immediately on the network, in contrast to the relatively slow dissemination in the physical world.
- One user can modify stored information, and other users on the network immediately see those modifications. For example, this can enhance the value of information as a tool for collaboration among users or organizations.
- The total information stored in the millions of computers on the Internet vastly exceeds what could fit in a personal computer's limited affordable storage capacity.

- The computers can process information in intensive ways, for example, to seek out more targeted or useful information.
- Computer-mediated information broadens the media from the printed page to include multimedia. The desktop computer's high-resolution screen is suitable for presentation, the computer can capture, store, and play back audio and video, and these media can be transported through the network. For example, the RealAudio and RealVideo media players allow audio and video content to be transported through the network and displayed on a desktop computer screen. A video media player is shown in Figure 2.5.

2.4.1 Finding Useful Information

Just capturing and storing large repositories of information doesn't directly meet a user's or organization's needs. Typically, an information repository is too large to be of value as a whole, but rather, users are interested in a targeted subset of the available information. Narrowing down information to a useful subset is a major function of information management, which the physical library addresses with information classification schemes, card catalogs, and reference librarians. Networked computing will eventually be much more effective at this, leveraging a computer's ability to process vast amounts of data.

User-Directed Access

The most straightforward way for a user to narrow down useful information is to question or interact with an application that has access to a large repository of information. There are three basic tactics listed in Table 2.7. The search exploits the computer's ability to systematically examine large volumes of data, if only the user can pose the right question. Browsing and navigating, on the other hand, allow the user to interactively guide the examination of the information, eventually honing in on useful information. The search is most useful in answering a specific question, whereas browsing and navigation are most useful when the user is curious or unsure of precisely what she seeks.

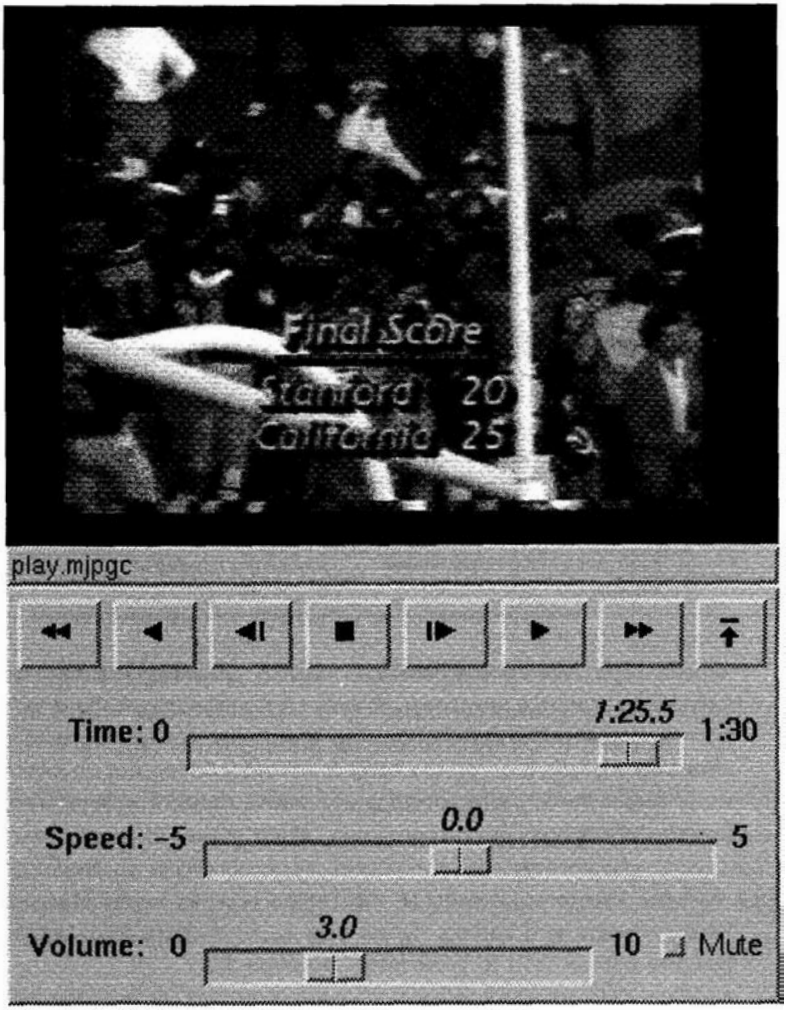

Figure 2.5 The continuous-media player allows viewing and control of a video/audio playback (from [Row92]). Source: the Berkeley Continuous Media Toolkit, a research project of the Berkeley Multimedia Research Center. Copyright © 1990–1998 by the Regents of the University of California. All rights reserved.

ANALOGY: An item search is like inquiring at the library about the availability of a specific book (title, author, and edition), and a topic search is like inquiring about the availability of any book on a

Table 2.7 Tactics for finding useful information.

Type of information retrieval	Definition	Example
Search	Pose a question, and information relevant to answering that question is returned. Item search: A narrow question with a precise answer. Topic search: A broader question about information on a broader topic.	Find articles in the medical literature containing the words "Parkinson's disease." The Web has search engines, such as AltaVista, that return pages containing specified keywords.
Browse	Examine a variety of repositories, hoping to opportunistically uncover interesting or useful information.	Follow a set of interesting Web hyperlinks. The Web hyperlink directly supports browsing (see the sidebar "World Wide Web" on page 37).
Navigate	Follow a map or similar navigational aid to arrive directly at the desired information and possibly related information.	Many Web sites have a site map that shows a complete tree of hyperlinks and subject headings.

specific topic (like cooking Mexican cuisine). Browsing is similar to wandering around a new town, at each corner heading in the direction that looks more interesting and promising. Navigating is like choosing a destination on the town map and working out an exact route to get there, perhaps noting alternative routes in case of traffic congestion.

EXAMPLE: *You may be familiar with some specific capabilities of Web browsers and servers supporting each of these four tactics (see the sidebar "World Wide Web" on page 37):*

- *Many Web sites specialize in answering queries, such as the latest selling price of a stock.*
- *The Web has a number of search engines that return Web content containing supplied keywords. Examples include AltaVista, HotBot, and Lycos.*
- *The Web hyperlink supports browsing. By adding hyperlinks, authors help users to find useful related information. The Web consists of a vast number of pages with hyperlinks among them.*
- *Web browsers make a history of accessed pages available on a menu, thus supporting a primitive form of navigation by allowing the user to backtrack to previously accessed pages.*

There is wide recognition that more sophisticated searching, browsing, and navigation mechanisms are needed, and this is the focus of research. Some techniques being explored include

- Searching will be based on an understanding of context of the information. Keyword searches may uncover information from distinct domains that can only be distinguished from the context (for example, "china" may apply to porcelain or a country).
- In the future, applications should automatically broaden a search to account for variations in terminology across organizations, fields, or nationalities, and may also perform automatic language translation. Simple keyword searches are complicated by different terminology used for the same concept (for example, a car "bonnet" in England is the same as a car "hood" in the United States).
- Browsing and navigation aids based on three-dimensional representations—exploiting users' ability to visualize information in three dimensions—will become commonplace.

Assistance from the Author or Publisher

The author and publisher can do much to assist the user in finding useful information, including several aids shown in Table 2.8.

As the world moves toward multimedia information, finding useful information is even more challenging. Indexing and metadata attached to audio, images, and video are particularly important, because inferring the content is difficult to automate (although this is being researched).

EXAMPLE: The user may want a picture of a brown Chihuahua from a database with images of dogs. With current technology, automatically distinguishing a picture of a Chihuahua from an Australian Kelpie is impractical (although it is more feasible to automatically recognize pictures of dogs). An index or metadata allows the dog's breed to be included in a search.

The idea of metadata is particularly useful for conveying other information about a document that requires human judgement and cannot easily be inferred from the document's content.

Table 2.8 Several ways an author can help the user find information.

Navigation aid	Description	Examples
Hyperlink	Link or reference from one document to another related document to assist in browsing and navigation.	Web hyperlinks lead browser to related pages (see the sidebar "World Wide Web" on page 37). References in this book lead you to related articles and books.
Index	A list of terms or subjects with hyperlinks or references to information about them.	Table of contents and index. Yahoo! provides an organized index to the Web.
Metadata	Think of metadata as "data about data." In the context of a document or unit of information, metadata is a description of the content (or possibly other attributes, such as quality, date of generation, etc.).	A short description of this book on the back cover summarizes its content. Other examples are the abstract of a paper, the title of a figure or table, or the textual description of image content.

EXAMPLE: The World Wide Web Consortium is standardizing a Platform for Internet Content Selection (PICS), which adds a form of metadata known as labels to Web pages [Res97b]. Unlike keyword searches, these labels can convey other attributes requiring human judgement, such as whether content is judged humorous, or offensive, or suitable for viewing by children.

Third-Party or Collective Recommendations

While assistance from an author or publisher of information is quite helpful, it does presume a relationship of confidence and trust between the publisher and user. It is useless in determining the authenticity or authority of information, or in conveying possibly unstated motives of the author or publisher (such as commercial gain, for example). Often, the judgement of other users or third-party authorities is helpful in overcoming these difficulties.

EXAMPLE: Consumers often depend on independent reviews of books, music CDs, and movies in a newspaper or magazine. They may rely on the collective judgement of other consumers, as represented, for example, by the appearance of a book or movie on a best-seller list. Or they may rely on informal word-of-mouth judge-

ments from their friends or the advice of others in a newsgroup or chatroom.

Information access over the network offers many opportunities to formalize and extend these mechanisms for identifying valued information. Computer-mediated systems that assist and augment the natural social process of exploiting a collective judgement are called recommender systems (or sometimes collaborative filtering) [Res97a]. In a sense, they are a hybrid between an information access and social application (as mentioned in "Recommendation Sharing" on page 38).

EXAMPLE: The on-line bookseller amazon.com provides recommendations on other books customers might consider buying (of course, in the interest of selling more books). One technique they use is to examine a database of books purchased by other customers. When the customer orders one book, amazon.com examines all the purchases of other customers who have purchased that same book. Written reviews submitted by other customers are also available.

Recommender systems are an area of research and commercial activity and will become increasingly common and sophisticated.

Third-Party Organization and Indexing: Digital Libraries

The Web is a large and growing repository of information, but it is also chaotic and disorganized. Its allows anybody to publish information that can be easily accessed by others, but its weaknesses include lack of control over what is published and the fact that there are many uncoordinated publishers. The Web supports browsing well, but search engines are of limited value because there is no structured representation of the data or organized indexing or metadata. Similar problems are addressed in a physical library by adding organization and indexing to a large body of published work emanating from many (otherwise uncoordinated) publishers. This is done by a third-party librarian, largely independent of the publication process.

If the structure and organization of a library are combined with representation and access in a networked computing infrastructure,

the result is called a digital library [Les97a, Les97b]. Subsets of the Web can be turned into digital libraries by adding indexing and navigation features.

EXAMPLE: Yahoo! is a rudimentary digital library that indexes a subset of the Web by subject. Many scholars maintain an organized index to the most important papers in their field and publish them for other scholars.

While some say that a vast collection of digital libraries will replace traditional paper-based libraries, this will take a very long time (if it happens at all). Many materials exist only on paper and would be time consuming and expensive to digitize. However, the digital library as a supplement to traditional libraries, particularly for newly minted information, is an idea whose time has come. For newly authored and published information, the digital library offers considerable cost savings by avoiding printing, physical distribution, and the geographic duplication of library materials. A digital library can also represent volatile and multimedia information and can be highly interactive.

There are several steps between author and information consumer, as listed in Table 2.9. None of these important functions evaporates with digital libraries—they all continue to add value. In digital form, they assume a different form, use a different medium, and may be more automated.

EXAMPLE: Various personal productivity applications assist in authoring, such as word processors, draw and paint programs, and music composition programs. Many print publishers are expanding to the Web as a publication medium. Yahoo! and others index the content of the Web. Many Web sites have site maps and search engines that provide navigational aids similar to (but far less capable than) the services of a reference librarian.

While technology can automate some of these functions, people will continue to add value. Keyword searches will identify pages relevant to a particular topic, but the difficult problem remains of information authenticity and reliability, which requires human judge-

Table *2.9* Steps from creation to consumption of information **[Sch95].**

Actor	Role	Examples
Author or performer	Creates information content.	Writes a book, performs a symphony, makes a music video.
Publisher	Verifies and improves quality, makes the work available for access or sale, controls use of its trademark.	Book publisher, record company, or movie studio.
Indexer	Classifies information and works.	Publisher of thesaurus, telephone yellow pages, library card catalog.
Librarian	Assists and guides user to appropriate content.	Reference librarian at the local library.

ment. A publisher also requires judgement to check an author's credentials, review information for accuracy, and invoke other quality-control measures. There is no complete substitute for a reference librarian, even for a digital library. The librarian can formulate more sophisticated searches after discussing user objectives and can often formulate better strategies for finding useful information.

2.4.2 Autonomous Information Sources

The user can determine what information is accessed, and when it is accessed, as discussed in the previous section. The broadcast is an extreme example of a diametrically opposite approach, in which the publisher determines what information is provided to the user and also when (see the section "Broadcasting" on page 36). The broadcast illustrates the extreme of an autonomous information source. In practice, many information access applications fall somewhere between user directed (sometimes called information pull) and autonomous source (sometimes called information push). There are different approaches, based primarily on three attributes of the application, as listed in Table 2.10. Most applications choose a mixture of these attributes and thus mix the pull and push models.

Table 2.10 Characteristics of information pull and push.

Characteristic	Pull extreme	Push extreme
Control	User requests specific, targeted information.	User subscribes to information on general topics.
Notification	User submits a specific standing question, which the publisher answers as appropriate.	Publisher provides appropriate notifications of useful information it thinks the user may want. The user can choose whether to access that information.
Timing	Information is provided at a time directed by the user (either immediately or at a scheduled time).	Publisher provides information at a time of its choosing. User may look at information as it is provided (direct style) or later (deferred style).

User Control: Subscriptions

The information-on-demand model is illustrated by the Web, where the user determines what pages to access and display. In the opposite subscription model, the publisher partitions information into what it believes to be natural categories or topics, each of which is called a channel. What information is actually delivered to each channel is under the control of the publisher. The user determines what channel(s) he or she wishes to receive and then makes a request (called a subscription) for the subsequent delivery of those channels.

ANALOGY: Broadcast radio and television, newspapers, and rnagazines—each representing a channel—are obtained by subscription. Each publisher differentiates the information content of its channels so as to attract the most subscribers.

EXAMPLE: PointCast is a push application that provides a standard set of channels on topics such as the stock price for a particular company, business news, or sports news. The user subscribes to the channels of interest, and thereafter they are presented on the screen as a screen saver. A screen of PointCast is illustrated in Figure 2.6, where (A) is where the user chooses the channel, (B) is a menu of articles for that channel, (C) is an advertising window (which supports this free service), and (D) is the article. The presen-

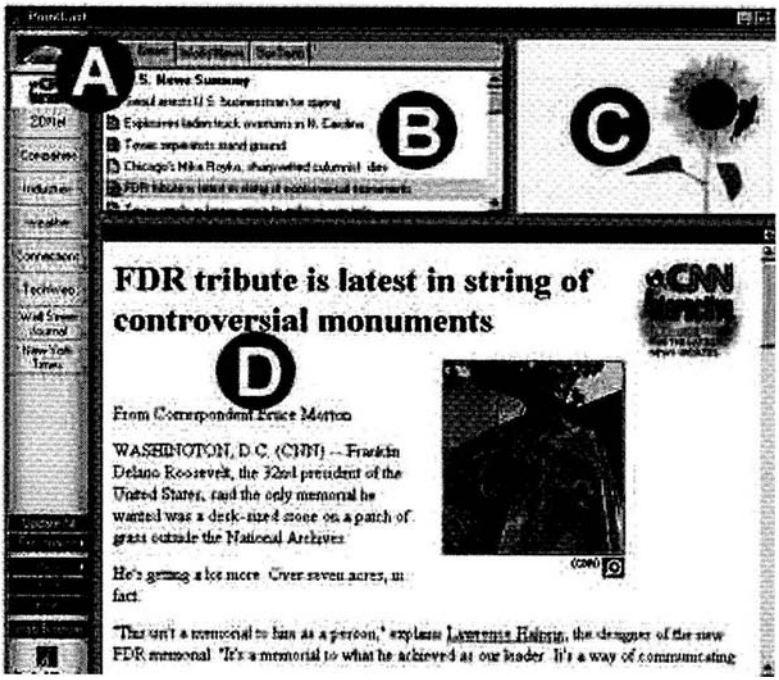

Figure 2.6 Illustration of Pointcast screen. Source: reprinted from PC Computing Online, August *1997.* Copyright © *1997* ZD, Inc.

tation is similar to a Web browser, except the provider chooses the content of each channel on behalf of the subscriber.

Subscriptions are particularly useful with volatile information, such as news reports or stock prices, since they remove the burden from the user of consciously requesting delivery. Further, if there is nothing interesting to report, the user is not bothered.

EXAMPLE: If a user subscribes to a channel carrying weather bulletins, information is provided only when the weather bureau issues a bulletin, so users do not have to remember to check periodically for bulletins. This is a coordination-monitoring application (see Table 2.5 on page 26).

User Awareness: Notifications

Rather than provide a stream of information, an alternative is to make users aware of what content is available and leave it up to them to determine what they wish to access. This approach is based on indexing the information in a channel and periodically pushing only the index—but not the indexed information—to the subscriber.

EXAMPLE: Many publishers of volatile information on the Web encourage users to subscribe to notifications of new content by email. These notifications alert users to new content, attracting their attention and giving them the option of accessing the content (using a Web browser, through a hyperlink embedded within the email message).

PointCast (see Figure 2.6) uses primarily a notification approach. The user is provided an index of stories available on each channel and chooses which ones to look at.

Timing

Users may or may not have control over the timing of subscribed information or notifications. If they do not have control, the information or notifications may still be stored for access at a time of their choosing.

EXAMPLE: PointCast (see Figure 2.6) controls when its information is provided to the user's screen, but the user chooses when to look at the stories.

An immediate form of information pull is appropriate for timely, volatile information, such as notification of a stock market crash or extreme weather event.

Push and Pull in Social Applications

The push/pull distinction arises in social as well as information access applications (see Section 2.3 on page 19), where the distinction is between one user initiating or responding to a request for interaction (or work group).

EXAMPLE: Email is a push application, because one user "pushes" a message at another. A newsgroup is a pull application, since a user must consciously go to the newsgroup to "pull" a message from

another user. In fact, the email and newsgroup applications can be viewed as push and pull variations of messaging.

In the context of social applications, each model has its strengths and weaknesses:

- Push is invasive, since a user never has total control over what is "pushed" at her. Excessive reliance on push can make a user's life unnecessarily cluttered and stressful.
- Pull requires a user to consciously initiate the interaction with another user or an information source. It does not alert the user to the need for an interaction and thus requires more attention and conscious activity on her part.

In most contexts it is best to combine push and pull (see the sidebar "Role of Push and Pull in Work Groups" for an example).

2.5 Education and Training

One promising networked application is education and training. Occasional predictions to the contrary notwithstanding, it is unlikely that networked computing will ever displace entirely the classroom or campus, as face-to-face interaction and socialization are important, particularly for younger students. It does offer considerable value as a supplement to the classroom and may replace it in some contexts (such as some corporate training and lifelong learning). It may also result in significant shifts in instructional techniques.

Several areas where networked applications offer obvious value include

- Remote learning: Networked computing empowers a student to work from home or office in a virtual classroom. This would benefit from all the remote conferencing capabilities and groupware discussed in Sections 2.3.3 and 2.3.4.
- Multimedia educational materials: A textbook equivalent offered on-line can use multimedia to supplement text and images. This offers not only direct value (for example, the ability to see the events of history rather than just read about them) but also excitement and interest (and dare we say entertainment?). This exploits the applications discussed in Section 2.4 on page 38.

Role of Push and Pull in Work Groups

Combining push and pull in a work group can minimize the burden on users while ensuring their appropriate involvement. Generally it can be said that

- Email (push) should be avoided for discussions and brainstorming, because it forces every group member to deal with every topic. Also, discussions are mixed with more time-critical interactions in the user's mailbox. The newsgroup allows each user to better control which topics she joins in as well as when she chooses to participate.
- Newsgroups and the Web (pull) cannot be relied on exclusively, because users can easily forget or ignore them. Users may be unaware that a topic of interest is being discussed and feel disenfranchised when they discover that they didn't participate in decisions.

The best modality for collaboration mixes the pull and push models:

- Any documents in a work group should be posted on the Web rather than sent in an email message (proprietary information can be protected using

(continued)

restricted access). This pull model allows users to control which documents to read at what time, and they always access the latest version.

- Brainstorming and discussions should also use pull, that is, a discussion forum application.
- The attention of users in the work group can be solicited by push, such as email. Typically, users should be informed by email when a new discussion topic is initiated or a new document is posted on the Web. Ideally, such informational messages should be sent only occasionally, aggregating discussions and documents.

As the number of people using social applications increases, using them properly becomes ever more critical. This is largely a matter of proper education of users.

- Just-in-time (JIT) training: For corporate workers, immediate applications such as the virtual classroom are increasingly ruled out by scheduling and travel constraints. At the same time, because of rapid change, deferred forms of education and training that empower people to learn opportunistically—at the time of need—are increasingly valuable.

EXAMPLE: DigitalThink is a successful start-up company providing training courses using the Web and also collaborative tools like chatrooms to allow students to interact with an instructor at a time of their choosing.

No doubt, radically new models of education and training will arise using combinations of immediate and deferred, and pull and push, applications. Education combines, in part, knowledge acquisition with the assistance of a teacher to offer assistance, explanation, and guidance. It is natural to combine networked access to self-learning materials together with tools for collaboration with a teacher and other students and trainees. This application area is ripe for innovation.

2.6 Business Applications

It is not an exaggeration to assert that networked computing is transforming the nature of business itself. There are numerous aspects to this transformation:

- Work groups: Social applications directly assist workers and managers in their jobs. Their immediate impact is in relaxing geographical constraints, allowing greater interaction with field personnel, and flattening the organizational hierarchy. They also increase the attractiveness of modern organizational techniques such as ad hoc task forces across functional departments, as they assist collaboration among busy, geographically dispersed workers.
- Operations: Business requires the movement of goods, information, and money within enterprises and among businesses. Increasingly, networked applications are replacing paper and the transportation system in the flow of information and money,

allowing better coordination (internally and externally) and reducing delay, administrative overhead, and errors.

- Electronic commerce: The transaction costs of firms dealing with one another and with consumers are being dramatically reduced, affecting firm boundaries and industry organization.
- Decision support: Managers and workers must make daily tactical and strategic decisions on what products to design and market, which suppliers to use, when and where to invest capital, etc. Networked computing can enhance the quantity and quality of information available to influence those decisions.
- Information management: Information and knowledge management applications allow enterprises to more widely distribute information internally and externally.
- Mass customization: The networked computing technologies enable greater customization of products and services sold to individual consumers.
- Consumer relationships: The rapid spread of the Internet enhances relationships between sellers and consumers, providing better customer support and new ways of selling.

All is not glorious, however. Networked computing as a means for cost cutting can make business relationships more impersonal and anonymous. It can automate the tedious and repetitive but can also eliminate low-skill jobs and reduce the autonomy of workers. Reliance on networked computing also introduces troubling vulnerabilities, as discussed in Chapter 8.

The evolution of these business applications parallels that of computing itself (see Section 1.1 on page 2) with four overlapping and coexisting phases listed in Table 2.1 1.

2.6.1 Departmental Applications

The typical enterprise is divided into specialized functional departments (accounting, manufacturing, purchasing, development, marketing, sales, etc.). The first business applications managed data within those departments, often replacing paper records.

Table 2.11 Four stages in the evolution of business applications.

Scope	Descriptive terms	Characteristics
Departmental	Client/server computing; enterprise databases; personal productivity applications	Move enterprise data from paper to electronic form, managed by mainframe databases.; add applications that exploit these databases and support workers, often on departmental servers; add decision-support functions and the ability of workers to access and manipulate data on their client PCs.
Enterprise	Business process reengineering; enterprise resource planning; data warehousing; data mining; groupware	Integration of databases is used to reengineer and automate business processes across the enterprise; data is consolidated for decision making.
Cross-enterprise	Electronic data interchange; electronic commerce; electronic business; supply-chain management	Data integration and business processes are extended to include other firms (both suppliers and customers); decision making is enhanced by real-time input; the Internet enables smaller companies to participate.
Consumers	Consumer electronic commerce	Provide shareholders and public with information about the company; provide consumers with marketing information; sell products and services and provide product support over the Internet.

EXAMPLE: Typical department-level applications would include human resources (tracking employees, doing payroll), accounting and finance, and managing manufacturing inventories.

A major application is on-line transaction processing (OLTP), which supports service agents who deal with customers, for example, to take orders, make reservations, take service or repair requests, or help customers make deposits or withdrawals. The large amount of data representing the information managed by these applications resides in a database management system (DBMS) (described further in Chapter 3). A DBMS excels at managing the type of data

common in business applications, which includes numbers (bank accounts, employee ages, etc.) and character strings (employee and customer names, etc.). Many data management functions in OLTP and other applications are commonplace across different business applications, and hence the DBMS has emerged as a separate software product category.

Departmental applications were originally developed and are managed by information systems (IS) organizations on mainframes. Since they are mission critical—the business cannot function without them—they must often run twenty-four hours a day, seven days a week (known as 24 by seven) and must be highly secure and reliable (see Chapter 8). The mainframe provides this secure, reliable centrally managed environment. On the other hand, such centralized systems and their associated IS organizations are, by themselves, rather inflexible [Wat95]:

- Workers and managers wanting customized information have to submit special requests, resulting in the infamous "IS backlog."
- Mainframes require large and expensive support staffs that tend to perpetuate the status quo.
- Since the IS organization serves the entire enterprise and imposes a uniform solution, individual departments have little ability to streamline or increase their own efficiency.

Decentralized computing addressed some of these issues by allowing workers to manipulate data in ad hoc and customized ways on their desktop personal computers (PCs)--especially using spreadsheets—without the assistance of an IS organization. The networking of PCs allowed workers to retrieve data from mainframe OLTP systems for this purpose. A primary factor in the success of the PC was its empowerment of individual workers, unleashing them from some constraints of the IS organization.

Decentralized computing also enabled individual departments to set up and operate applications themselves, based on relatively inexpensive microprocessor-based computers called servers. This led to client/server computing at the departmental level, where workers' PCs (called clients) could store, access, and manipulate

data stored in departmental servers or mainframes. Installing and managing applications within the departments added flexibility, and local control resulted in greater efficiency and streamlining. Client/server computing is discussed further in Chapter 3.

Client/serer computing unfortunately also allows a proliferation of incompatible systems and applications across the enterprise. The computing environment became more heterogeneous, increasing maintenance and operational difficulties, and placing obstacles in the way of enterprise applications, as discussed next.

2.6.2 Enterprise Applications

The bottom line of the enterprise—the production and sales of products and services—requires the cooperation and coordination of departments across an enterprise.

EXAMPLE: Manufacturing requires a flow of parts from suppliers, which in turn requires purchasing to order the parts and accounting to pay suppliers. Further, manufacturing must coordinate with sales to forecast demand and maintain manageable inventories, which in turn affects purchasing.

Organizational theory defines the business process as a stream of related activities starting with the acquisition of resources (people, capital, supplies, etc.) and resulting in a product or service in the hands of customers [Dav93]. A key term in this definition is stream, which indicates that the constituent activities are repetitive. This distinguishes them from the social and information management applications discussed earlier, which tend to be ad hoc: Each interaction or collaboration is different.

The structure of a business process is illustrated in Figure 2.7. It takes resources and services from suppliers and ultimately provides products or services to customers. The process requires coordinated activities within different functional departments, as well as the flow of material, finished goods, and information among them. It focuses on the activities to be coordinated and the pattern of information and material flows. The design of a business process

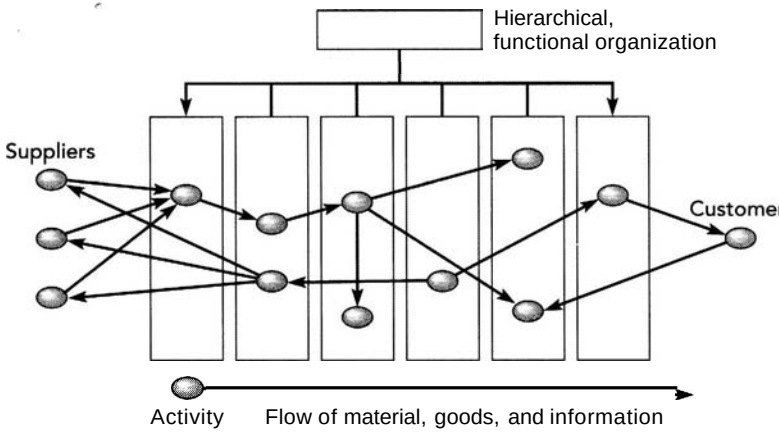

Figure 2.7 Illustration of a business process, which horizontally spans a typical vertical functional organization.

asks how the operation of the business can be made more efficient and effective by somehow modifying both activities and flows, finding a better partitioning of functions among workers and networked computing, and a better interface between workers and technology.

While networked computing cannot contribute directly to the flow of people, material, or finished goods in a business process, it is ideal for controlling that flow and the manipulation of information supporting its coordination. A business process application is similar to social applications in that it enables the communication among and coordination of departmental organizations, instead of users (see Section 2.3.2 on page 22).

EXAMPLE: In manufacturing, networked computing can support the flow of requests for material and parts from manufacturing to purchasing (so that suppliers can be notified), from receiving to accounting (so that suppliers can be paid), and from sales and distribution to manufacturing (forecasting the volume of production needed), among others. It can also coordinate their actions, such as matching manufacturing volume to sales forecasts.

Workflow

Where an application primarily supports a group of workers cooperating on repetitive tasks, a business process is often called workflow. It is another class of social application, in which the tasks are ongoing and repetitive.

EXAMPLE: In a customer service application, telephone agents typically interact with customers to ascertain their problems, and support technicians solve them. Each validated problem results in a trouble ticket, which flows from telephone agent to a group of support technicians. Within the group, each trouble ticket is routed to an appropriate and available technician. The solution is reported to the telephone agents to inform the customer.

Business Process Reengineering

Because a business process involves repetitive activities, it can be planned and optimized in advance. This is termed the engineering of the business process, the goals of which include

- Minimize operational cost (personnel costs, inventory and transaction costs, etc.).
- Reduce delays required from inputs to customer, allowing the process to respond more quickly to customers and to changes in the marketplace.
- Better quantify and minimize the costs associated with the overall process (as opposed to costs at the department level).
- Provide improved and more timely information to support business decisions.
- Make the process flexible, so products can easily be changed in response to marketplace needs, or even customized to individual customer's needs (called mass customization).

Used effectively, networked computing can provide considerable value in each of these areas. However, most existing processes were designed years ago—or arose in an incremental and ad hoc manner—and fail to make the most effective use of computing.

It is popular to take a fresh approach to the design of a business process in a way that uses networked computing effectively. Called business process reengineering (BPR) [Dav93], where done effectively, this has resulted in major savings and competitive advantage. (The term BPR has lost some of its appeal because of its association with mass layoffs. A term used more recently is business transformation.) Business process reengineering is not only about networked computing but also accounts for all aspects of the process, including the contribution and organization of workers and the flow of material, supplies, and finished goods. An essential issue in the design of a business process is what people do particularly well, what networked computing does particularly well, and how the two can work effectively together. Simply automating an existing process is rarely the most effective use of networked computing (see the sidebar "Electrification: Lessons from an Earlier Technological Advance" on page 8 for a useful analogy).

The phases of a BPR project consist of

- An analysis of business requirements and costs.
- A design that structures the individual activities and the flow of materials and information among these activities.
- A development of the management and worker organizational systems, as well as computer systems and software.
- The deployment of the application, which includes worker training and installing and testing the information systems.
- The operation of the application supporting the production, sales, and distribution of goods and services.

The design phase brings managers (who understand the business goals and the organizational and human resource challenges) and technologists (who understand the capabilities and limitations of the information technologies) together, but it benefits greatly from participants who understand both the management and the technology challenges. This book will prepare you to participate in this. The design and deployment of new applications is discussed further in Chapter 3.

SAP: Largest ERP Vendor

SAP (Systems, Applications, and Products in Data Processing), headquarteredin Walldorf, Germany, is the largest ERP vendor. Others are PeopleSoft, Baan, and Oracle.

SAPS most recent product offering is called R/3. It incorporates many of the software technologies described in Chapter 6 (frameworks, components, cross-platform object communication). SAP's most successful products support manufacturing, accounting, logistics, and human resources. In 1998 it is in the process of expanding its applications to the extended enterprise with supply-chain management.

Enterprise Resource Planning

The computing component of many business processes was designed and implemented by an IS department working with business units, often utilizing the services of outside consultants and professional services companies. Some of this custom design and implementation is unnecessary because — while there are always local differences — there is much commonality among the processes in different companies.

EXAMPLE: Accounting systems have a common purpose and operate under the same accounting rules. Human resource management (payroll, benefits, taxes, etc.) is reasonably standardized across companies.

It may be reasonable to purchase, rather than develop from scratch, an application to support a given business process. Enterprise resource planning (ERP) applications supply a major piece of an application, with sophisticated configuration tools to customize to local needs. (The software methodologies making this possible are described further in Chapter 6.) ERP products are available for business processes that are fairly standardized across different companies, such as the following:

Sales-force automation supports the field sales force with automated order entry, availability, and delivery and pricing information, for example.

- Document management supports the massive documentation of designs, parts catalogs, etc.

Customer service and support coordinates telephone agents and technicians supporting customers.

- Manufacturing logistics provides logistical support for the flow of materials and finished goods, inventory control, purchasing and payments to suppliers, etc.
- Accounting allows the company to keep track of cash flow, assets, profits and losses, etc.

- Human resources tracks employee history, salary and benefits, tax payments, etc.
- Supply-chain management coordinates with suppliers and customers, for example, in the supply relationships that support a finished-goods manufacturer.

A key to ERP is data integration—meaning that data associated with the application in question and previously closely held within departments, must be integrated into a single application. (Given the heterogeneity of computers and software in most enterprises, this is a severe challenge, discussed further in Chapter 3.) In effect, an ERP application serves as a framework that allows information to flow automatically among department-level systems. ERP also provides several related functions:

- Forecast and planning: Effective operations dictate that many resources (such as real estate and worker head count) be planned well in advance, and suppliers may need considerable lead time to manufacture and ship needed supplies. The ERP application can support these functions by providing historical data and future projections.
- Control: Efficiency can be improved by actively controlling the flow of goods to minimize inventory and manage accounts receivable and cash. ERP can provide feedback through the chain of activities in the process to control and coordinate them.
- Real-time monitoring: The ERP application can observe the business process as it is happening and present summary performance information. It can also highlight problems—such as weather-related transportation delays—allowing managers to react.

The downside of ERP applications is that they require an enterprise to adapt its processes to the software, whereas ideally it would be the other way around. Also, planning and deployment typically utilize experienced but expensive consultants (often accounting firms that have entered this business).

Integrating Multiple Applications

As discussed in Section 1.2, a networked application is frequently embedded in a large-scale system, of which networked computing is only an element. Business applications illustrate this well, since any business process involves workers and materials, and often customers and suppliers, as well as computing systems. Frequently, more than one application is involved in a business process, and any given department has to deal with multiple processes and applications.

EXAMPLE: A manufacturing process necessarily involves financial, inventory management, supply-chain, and shipping processes, among others. The manufacturing department is affected by all these processes, and data residing there may be manipulated by them all.

Decision Support

ERP forecasting and monitoring illustrate the important role of networked computing in not only operations but also support of managerial decisions. Networked computing offers capabilities of particular benefit to this decision support:

- Timely information: The networking portion of networked computing communicates data around the world with negligible delay (seconds or fractions of a second). As a result, data that supports decisions is available almost "as it happens."
- Data reduction and presentation: In a typical business application, far too much data flows for a person to absorb in raw form. The computing portion of networked computing allows this raw data to be aggregated and summarized.
- Knowledge acquisition and management: Computers are increasingly capable of aiding workers in acquiring knowledge from massive amounts of organized information.

Data warehouses and data mining are applications directed at decision support (see the sidebar "Data Warehouses and Data Mining").

Data Warehouses and Data Mining

Data warehousing and data mining are two applications tailored to decision support, addressing some limitations of OLTP applications. Decisions should take into account not only where things are now but also how they got that way, and use an enterprise-wide picture. A data warehouse is a very large nonoperational database (managed by a DBMS) that systematically captures information from a number of operational OLTP databases. It provides two major decision-support benefits:

- History: OLTP systems represent the present, but the day-to-day operational needs usually do not require maintaining a well-structured operational history. A data warehouse systematically captures the past. It is not updated on a transaction-by-transaction basis as in OLTP, but rather by periodically capturing snapshots of data from OLTP databases.
- Consolidation: OLTP applications typically don't present the total picture, because relevant data may be spread across major databases. Day-to-day operational needs may require only information for a single department or

Knowledge Management

A working definition of knowledge in computing is large amounts of information suitably organized and indexed to make it more useful to people. Knowledge is especially crucial to good decision making and is also a key competitive asset of a modern enterprise. Historically, the acquisition and management of this knowledge was very informal. Most knowledge resided in the heads of managers and workers, or in large repositories of paper documents, and knowledge was shared among employees on an ad hoc basis. Modern trends such as downsizing and layoffs, higher turnover, and geographical dispersion of workers point to the need for more formal and structured approaches to knowledge management [Ole98]. Networked computing can aid the acquisition of knowledge by allowing workers access to large information repositories that can be systematically searched or navigated, and, with the aid of the network, knowledge can be widely dispersed throughout an enterprise. Recognizing the increasing importance of knowledge management, by 1998, 40 percent of the Fortune 1,000 corporations had put knowledge management on a similar plane with the management of technology, money, real estate, and other corporate assets by creating a "chief knowledge officer" responsible for creating the infrastructure and culture of knowledge sharing and management [Ole98].

EXAMPLE: Consider a customer service organization in which service agents encounter and solve customer problems and enter this into a knowledge base. This acquired knowledge is immediately available to all service agents, increasing responsiveness and efficiency in solving similar problems. This knowledge—made immediately available to the manufacturing and product development—can be used to discover flaws in manufacturing processes or to improve designs.

Knowledge warehouses are databases (again managed by a DBMS) storing qualitative rather than quantitative information—work manuals, documentation, proposals, employee information, directories, newsletters, etc.-previously stored on paper. Increasingly, the Web browser is used as the interface to a knowledge base, allowing knowledge to be geographically dispersed across different

a single business process. A data warehouse consolidates information from multiple databases to gain an overall picture of the business operation.

Overall, the goal of a data warehouse is to present a consistent and correct historical image of an entire business, or at least a business process.

Data mining is an application that looks for unexpected patterns in large amounts of data. Unlike database queries, data mining does not require the user to ask a question, but rather tries to identify what questions should have been asked. Its emphasis is predicting future trends in a business by uncovering patterns within the massive data in a data warehouse. In a limited way, data mining attempts to extract knowledge from the information residing in a business process.

EXAMPLE: Data mining is put to a number of uses [OHE96a], such as finding which medical treatments are most effective, uncovering the relationship of personal characteristics to voting patterns for courtroom jurists, relating credit card customer characteristics to the likelihood of default, finding risky behavior patterns for insurance companies, and performing technical analysis of stock prices.

repositories and made readily accessible to managers and workers. The knowledge warehouse is essentially a digital library applied to a business application (see "Third-Party Organization and Indexing: Digital Libraries" on page 45).

2.6.3 Cross-Enterprise Applications: Electronic Commerce

Once business processes within an enterprise employ networked computing, and decision-support tools are in place, there remains a major opportunity. A primary operational function of any enterprise—selling goods and services to other enterprises and to individual consumers—can also exploit networked computing, particularly with the growing ubiquity of the Internet. The selling of goods and services among enterprises comes under the general heading of commerce, and when supported by networked applications, it is called electronic commerce. Electronic commerce provides many of the same benefits for groups of enterprises and individual consumers as business process reengineering within the enterprise, including reducing costs (administrative and overhead), delay, and errors. Electronic commerce also opens up entirely new channels, such as selling goods and services directly to consumers over the internet, and improves management decisions by providing more timely, complete, and accurate information [Kee97].

Any commercial transaction involves three basic steps shown in Table 2.12. Each of these steps—if completed over the network-presents its own challenges. The matching of buyers and sellers is an information management challenge, the negotiation of terms is a social application, and the consummation raises issues of how payments are rendered and information goods delivered securely over the network.

Electronic commerce invokes many similar operational needs as enterprise applications but presents many more constraints and obstacles:

- Firms must maintain a hands-off business relationship, and thus the information flowing must be restricted and controlled. Inad-

Table 2.1 2 Three steps in a typical electronic commerce transaction.

Stage	Description	Typical mechanisms
Matching buyers and sellers	The seller makes the buyer aware of what is available for sale. This is largely an information management problem (see Section 2.4 on page 38).	The seller has to make the buyer aware of goods available, through mechanisms such as advertising, on-line catalog, recommender system (see "Third-Party or Collective Recommendations" on page 44).
Negotiating terms	The buyer and seller reach terms and conditions on the sale, including price, delivery schedule, etc.	Negotiations are supported by work group applications (see Section 2.3.4 on page 28). If multiple buyers and sellers are involved, this may include an auction.
Consummation	The agreed sale is completed by transfer of goods and payment.	Key steps are order: buyer places order with seller (in a form that cannot be repudiated); fulfillment: seller conveys goods to buyer (in a form that cannot be denied); escrow: ensure delivery of goods by putting payment in escrow with a trusted third party; payment: buyer pays seller for goods (in a form that can be proven).

vertent disclosure of proprietary information is a serious security breach.

- Commercial relationships always have associated payments, which is often not a requirement of an enterprise application.
- Commercial relationships are often not conducted at fixed pre-agreed prices, but there may be a market involving competitive bidding, negotiation of terms, etc.

Particularly interesting are the possibilities for real-time marketplaces that work similarly to the stock market, something that is not too practical in ordinary retailing.

EXAMPLE: The Fastparts Trading Exchange is a Web-based electronic marketplace that allows manufacturers, contract assemblers,

component manufacturers, and franchised distributors to sell electronic components to one another. Companies with surplus components offer them for sale in an electronic marketplace with real-time negotiation. FastParts handles the payment and shipping of the parts and certifies sellers to ensure quality. FastParts ensures the negotiation and transaction are anonymous, so companies cannot communicate competitive information.

eBay is an on-line auction targeted at individuals. Contrary to FastParts, it is completely public. Sellers and bidders are known to each other. Each bidder typically negotiates terms (delivery, payment method, etc.) with the seller before bidding and can inquire about the goods offered for sale by email. Buyers and sellers can rate one another, and these ratings and other details of each auction are available on the site.

Cross-enterprise applications also introduce administrative and technical challenges, because coordination and control is difficult across administrative boundaries. While enterprises can form bilateral coordination agreements with suppliers and customers, each supplier has its own customers and suppliers desiring similar coordination. The market is a complicated web of suppliers and customers, and since each enterprise can't afford a proliferation of systems, achieving the full potential of cross-enterprise applications requires that all enterprises adopt compatible solutions. This is a major technical challenge, especially in light of heterogeneous legacy systems.

EXAMPLE: The U.S. automobile companies (General Motors, Ford, and Chrysler) have been among the most aggressive in adopting electronic commerce. They have formed the Automotive Network Exchange (ANX) connecting the automobile companies and suppliers. Members can use the network for electronic commerce even if the automobile companies are not involved. The reach of the automobile company suppliers is great, so ANX could form the basis for a much broader electronic commerce interest group [Jon97].

The elements of electronic commerce listed in Table 2.12 have come together gradually and incrementally—only recently have commercial transactions been conducted in totality over the network. Cross-enterprise electronic commerce has passed through

Legacy Applications and the Year 2000 Problem

Once applications are installed, they're often around a long time. Illustrating this is the Year 2000 (Y2K) problem. Many programmers in the 1950s and 1960s didn't anticipate that their applications would still be operational forty years later, and they saved memory and storage (precious commodities at the time) by truncating years to the last two digits (assuming the first two are always 19). As "January 1, 2000" is misconstrued as "January 1, 1900," havoc may result. These programs are written in old computer languages, and the original programmers have long since retired, so it is difficult to find and repair all instances.

The extent of Y2K points to the surprising number of legacy applications—those using obsolete technology but still in operation. Replacing a legacy application is time consuming and expensive and involves difficult logistical challenges. More sobering is the thought that new applications deployed today may have long operational lifetimes. On the positive side, Y2K has itself stimulated many replacements of legacy applications. Bypassing Y2K is one by-product of installing an ERP application, for example.

three increasingly ambitious phases: electronic data interchange, electronic money management, and electronic business logistics [Kee97]. Each of these phases expanded the role of networked computing. Technical challenges in business-to-consumer electronic commerce are discussed in Chapters 8 and 10 (see "Consumer Electronic Commerce" on page 38).

Electronic Data Interchange

Electronic data interchange (EDI)-the exchange of business messages in industry-specific standard forms—is a longtime electronic commerce application. EDI replaces paper supply/customer documentation, such as purchase orders and invoices, with electronic messages. Although it has been available for many years, it is usually associated with private communication links between the largest firms. The public Internet and the Web make EDI accessible to small companies. In addition, the interactive Web allows EDI to transcend computer-to-computer links and allow workers in one company to access supporting systems in another company.

EXAMPLE: Boeing's Part Analysis and Requirements Tracking (PART) system uses the Web to extend EDI to their small-airline customers. It streamlines Boeing's organization by allowing customers to perform searches for obscure parts without help from Boeing workers. Significantly, PART also allows an airline's mechanics and other frontline workers to directly search for and order parts without help from purchasing.

Electronic Money Management

Terms of sale must be negotiated and goods and services must be paid for. The second phase of electronic commerce—addressing the payments—is electronic money management, which is evolving from the province of large companies exchanging payments using banks as intermediaries (called financial EDI, or FEDI), to empowering smaller firms and even individual consumers to electronically manage payments. FEDI authorizes electronic funds transfers between bank accounts.

A major goal is the reduction of transaction costs, making electronic money management practical for smaller payments among small firms and from individual consumers. Even payments of a thou-

Dell Computer and Mass Customization

Dell Computer has eliminated traditional distribution channels, thus reducing costs in direct and indirect ways. Customers access the Dell Web site to view a product catalog and custom configure a computer to their budget and requirements and obtain an immediate price quote. Once an order is input, it generates immediate feedback to various internal business processes—such as manufacturing and purchasing—and directly to parts suppliers. Illustrating mass customization, a computer is manufactured to specification after it is ordered, rather than retrofitted by a retailer, distributor, or customer.

sandth of a cent (called a millicent) may be needed for selling individual snippets of information on the Web. While FEDI deals strictly in the transfer of money between bank accounts, credit card payments and cash allow direct payment from buyer to seller (although credit cards require later involvement of financial institutions). There needs to be the cyberspace equivalent of these payment mechanisms. However, where individual consumers and the Internet are involved, security and privacy are major concerns: Consumers are concerned that credit card numbers may be stolen or purchases may be tracked.

EXAMPLE: Secure Electronic Transactions (SET)—an initiative of Visa International and Mastercard—enables consumers to more securely make credit card purchases over the network or in person. SET ensures the consumer is the legitimate card holder (called authentication), precludes the merchant selling the goods from seeing credit card numbers or other financial information, and precludes financial institutions from tracking purchases.

Chapter 8 discusses secure electronic payments suitable for individual consumers, including credit card transactions and digital cash (the electronic equivalent of the cash you carry in your wallet).

Business Logistics

The third and final stage of electronic commerce is integrated business logistics [Kee97]. Logistics encompasses a broad swath of processes—transcending buying and selling—including coordination of material flows, finished goods, services, and people. It obtains efficiency, speed, accuracy improvements, and cost savings, not simply by automating existing paper-based processes, but by changing the very nature of business—reorganizing it around networked computing technologies and particularly the public Internet. An example of the possibilities is mass customization (see the sidebar "Dell Computer and Mass Customization"). Logistics extends the business process to the business-to-business processes in the customer-supplier relationship, reducing organizational barriers and flattening organizations, lowering barriers to those customers and suppliers, and integrating internal operations with suppliers and business partners. Logistics is heightening customer expectations for speed, accuracy, customization, and cost.

EXAMPLE: Dell's business model has many advantages. Manufacturing immediately builds and ships a customer's computer, requiring no inventory of finished goods. A procurement process can respond quickly by transmitting orders to suppliers, minimizing inventory costs. When new pricing or product strategies are introduced, marketing obtains immediate feedback on their success ***[SAP97].***

ERP vendors are embracing supply-chain management (see "Enterprise Resource Planning" on page 60). This class of cross-enterprise application extends ERP to controlling and monitoring the flow of materials, goods, services, and money through a supply chain. In its most extreme form, supply-chain management can create virtual enterprises [Gre96]. In this model of business, firms aren't tempted to accumulate all the specialties necessary for the manufacturing and marketing of products. Each individual product is designed, manufactured, and marketed by opportunistically pulling together (and later disbanding) an alliance of firms. In theory, virtual enterprises can react more quickly to market opportunities, and competitively chosen partners can deliver high quality, responsiveness, and low cost (compared to an internal supplier). The impact of networked computing on the boundaries among firms is discussed further in Chapter 5.

2.6.4 Consumer Applications

Consumer applications allow enterprises to interact with individual shareholders and consumers and provide a major new channel for marketing and selling goods and services to consumers. Unlike cross-enterprise applications, where proprietary data networks are an option, before the Internet, it was difficult to reach consumers electronically for marketing, sales, and post-sale support. Such capabilities have been available using the telephone for many years—for example, catalog sales—but the Internet offers a richer interaction, the tailoring of sales techniques and pricing to individual consumers (see the sidebar "amazon.com: On-Line Merchant" for an example), and lower transaction costs.

amazon.com: On-Line Merchant

A successful merchant that molded itself around consumer electronic commerce is the bookseller arnazon.com (even its name reflects an exclusive on-line presence). Like any retailer, amazon.com is an intermediary between distributors and customers, but since its customers browse on-line, amazon.com need not maintain an inventory. Rather, it passes orders directly to distributors, in an illustration of supply-chain management.

Amazon.com also illustrates that on-line retailing can offer new features and services. For example, each book the customer considers is accompanied by reviews from other customers (a recommender system), and additional sales are encouraged by listing related titles (based on both topic sales and customer behavior). The interests and behavior of customers can be monitored and presentations tailored accordingly. The customer can request email notification of books in categories of interest (an information push application).

2.7 Similarity of Social Systems and Networked Computing

This chapter has emphasized that the unique applications of networked computing are intimately interwoven with social systems of all types: the citizenry, interest groups, and organizations pursuing business, education, and commerce. Although probably not evident yet, there are remarkable similarities between the types of issues addressed in this chapter — how networked applications serve users and organizations — and how networked applications are internally organized and implemented, which is the subject of much of the remainder of this book. Examples of this abound:

- Organizations and other social systems must delegate responsibilities to different people, and networked applications must delegate responsibilities to different computers.
- People must communicate and coordinate themselves to accomplish defined tasks and so must the pieces of a networked application.
- People must figure out how to avoid conflicts when they share a common task or resource and so must networked computers.
- Society must allocate finite resources efficiently — such as highways, water, electricity, etc.-and so must the network and the computers attached to it.
- Social systems must cope with breakdowns, natural disasters, criminal behavior, and such occurrences, and so must computer systems.

You are urged to reflect on these parallels throughout the remainder of the book. Many of the same issues arise in social contexts and in networked applications, and many of the solutions are directly analogous, or at least recognizable. This should help you navigate the many technical issues to be uncovered subsequently, and the book aids this by providing numerous analogies. The ideas and concepts applied to the design of computer systems may also suggest similar concepts that can be exploited in the design of social systems, such as business organizations.

2.8 Open Issues

Applications raise a number of open issues: How are they invented? What impact do they have?

2.8.1 The Productivity Quandary

Economists find it difficult to discern productivity gains attributable to networked computing. Macroscopically, in the United States overall productivity growth has slowed the last couple of decades just as networked computing became more prevalent. Microscopically, correlations between the financial results of individual firms and networked computing investments are difficult to discern. While these circumstances seem counterintuitive, there are a few relevant observations:

- Networked computing has its biggest impact on the service sector, where productivity is notoriously difficult to measure.
- Much of the benefit is in quality, which is difficult to separate from productivity when measuring improvement. For example, networked computing empowers solitary users to produce higher-quality documents (but perhaps consuming more time) and run financial projections not previously possible (but also consuming time).
- Most importantly, automating existing processes is not where the major productivity gains are expected. Rather, both the processes and the social dimension of the organizations that incorporate them have to be reengineered to make most effective use of the technology [Wa189] (see the sidebar "Electrification: Lessons from an Earlier Technological Advance" on page 8 for a historical analogy).

So, the question is, "When will substantial and documented productivity gains from networked computing technologies appear?" Only time will tell.

2.8.2 How Are New Business Applications Invented and Developed?

Automation of existing processes does not fully realize the potential of networked computing. The biggest gains are achieved when application design is integral to the overall process design, including human resource and organizational design. Application design is thus not primarily a technical activity, and yet the technical issues of performance, cost-effectiveness, and flexibility are also important. Thus, the traditional gulf between nontechnical and technical workers, where the former rely on the latter to define the applications and infrastructure and the latter rely on the former to define the application context, cannot meet the challenge of seamless interworking of the organization and its enabling networked applications. Can a collaboration of management with information technologists continue to work, or is a new cadre of technically astute managers (or managerially astute technicians) needed?

2.8.3 The Glut of Information and Communications

Both communication and information access become so easy that many users find their biggest problem is too *much* information and too *much* communication. Networked computing can transmit and manipulate astounding amounts of information and thus can filter, search, and present it in helpful ways. Recommender systems can also be very helpful. Communications and collaboration are arguably more problematic, since human communication is inherently "push," performed invasively and consciously [Mes96b]. As the global networked citizenry expands, the demands of communication expand with it, consuming more of the user's time. In response — to avoid the constant interruptions and scheduling problems — more interaction becomes deferred (email, voicemail, etc.), which is less invasive but also less efficient for many purposes.

Increasingly, applications must incorporate sophisticated capabilities to improve the quality and limit the quantity of information, communications, etc. Nevertheless, the growth of information and communications seems to be exceeding the capability of these technological measures. Certainly more research and commercial activity are needed in this area.

2.8.4 Accommodating Change

An obvious attribute of modern business is continual change. Firms recast their businesses rapidly, the product birth and obsolescence cycle shortens, reorganization is endemic, and networked computing penetrates more deeply. The boundaries of firms change with relentless mergers, acquisitions, and divestitures. These environmental factors place severe stress on information systems, which would ideally ease the path to change rather than inhibit it. An important objective—but one not achieved with the current state of technology—is flexible applications that readily adapt to changing needs. This is a serious challenge, not only to the applications themselves but also to the infrastructure that supports them, and one that is arguably not met by today's business applications.

Further Reading

There are not yet other books covering social applications of networked computing that can be recommended. [Les97b] is a good treatment of information management (especially digital libraries), [Ole98] is a concise and readable introduction to knowledge management, and [Kee97] is a good introduction to electronic commerce, with an extensive glossary.