can have on the user's experience and how user requirements drive systems design. In this section we would like to present some applications that illustrate why we should go to all this trouble and what people gain from using CVEs.

Generally, there is an impetus for considering the use of CVEs for work. As noted above, many studies have pointed out many examples of co-located work that is achieved by the participants communicating with each other, and able to look at and manipulate some mutually available workplace artefact. These artefacts can be as diverse as timetables, displays of train lines or of stock movements in the stock exchange, or even parts of a person who is being operated on by diverse medical and nursing professionals. As Chapter 7 points out, the promise of a CVE is the ability to manipulate these virtual objects and share them with remote partners.

We may, then, be able to utilize and share "physical" or, as Chapter 7 points out, *visual* elements of our work – such diverse things as architectural visualizations, whiteboards, X-ray slides, video images and even "3D objects and animations intended for the 3D world itself". In each case, there are supporting arguments for the benefits of the ability to collaborate around these artefacts while geographically remote.

However, this may not be all we can do. If we incorporate physical objects in some way into the CVE we can actually manipulate them. Such an object-focused element in a CVE promises much. An example of an early version of such a CVE is that discussed in Chapter 11. This chapter describes how a hybrid CVE/media-space system can allow scientists to collaborate in ways similar to how they would naturally collaborate around rare pieces of equipment in the real world. Given that a particular country might only have a very small number (if any) of certain types of exotic equipment, this illustrates how CVE systems can fulfil a real need. Continuing on the scientific theme, Chapter 12 illustrates how CVEs can be used for scientific visualization and education and Chapter 10 shows a 3D interface to support Web browsing.

1.3.5 Emerging and Existing Virtual Cultures

While it may be a truism that, deep down, people are all the same, it is also true that people act differently in different situations. A typical example is how a person will behave at home, at work and at a football match. Each situation lends itself to a different behaviour. Also, each situation will often be experienced with different people. People's behaviour may (will) evolve over time as they grow accustomed to a situation and group of people – for example the newcomer in the office may be quieter than normal until he or she learns how to get on with their new colleagues. As we have noted above, social activity depends on place, people and time and combinations of these will gradually evolve a set of (normally) unspoken and implicit rules that demarcate regions for activity and govern conduct (Giddens, 1984; Harrison and Dourish, 1996; Munro, 1999; de Certeau *et al.*, 1998).

Since CVEs construct places and involve multiple people, it is perhaps not surprising to see customs and cultures emerge that are built around virtual places. For a simple proof of this, one only has to sample some of the mailing lists and

USENET newsgroups available on the Internet – each group has its own customs and accepted behaviours (e.g. Mynatt *et al.*, 1998, 1999; Preece, 1999). Given the richness of a CVE (text or graphical) compared to a mailing list or newsgroup we should expect to see cultures form in long-lived CVEs, and this indeed proves to be the case. Chapter 14 illustrates how a work-related CVE supports social interaction that in turn supports collaborative work. Chapter 15 describes how a virtual office allows the author to make himself available to others in ways not previously possible. Finally, Chapter 13 describes an experiment in which people are assigned specific roles in a virtual culture and given the chance to experience life from a viewpoint that might normally be unavailable to them.

1.4 Summary

In this chapter we have tried to provide a basic introduction to some of the issues that we consider are most important for CVEs at the current time. We have not tried to discuss virtual environments in general, since there are already many good books on the subject. Instead, we have focused on the challenges required to go from a single user environment to one that can actually support rich and meaningful communication and collaboration between groups of people. We have divided the book into a number of sections, each of which represents a particular theme, such as representations of self and other, understanding the viewpoints of others, technical challenges, applications and emerging cultures in CVEs. We hope that the following chapters prove useful and informative.

system do the walking". The user's responsibility for controlling their avatar is removed and made a task of the system. In their general contact space, and in their meeting space, users do not have close control of their avatars. Rather, if they wish to do certain things, like leave, the system makes the avatar go through a series of stereotyped preparatory behaviours, such as gathering papers together, lining them up on the desk, and then getting up.

In allowing the system this control through scripting of stereotyped behaviours, it takes the onus off the user for fine-grained control. Thus it should be possible, the authors suggest, to access such a system from both a high-end machine and also a lower end machine with a modem, or even some kind of mobile device.

This solution has a number of implications: it allows a certain freedom of the user from the basics of interaction, but creates a number of burdens on the design of the repertoire of behaviours. The behavioural repertoire will have to be flexible, or limited to a small number of domain-specific behaviours, which could limit the system's usefulness. At the moment, this system is not yet in serious use, so it will be interesting to see what issues come up. Would such a solution even work in a more "spatial" CVE like MASSIVE, where the user is required to do a more spatial task? How would one begin to design a putative system that had the sort of repertoire of behaviours needed for a more spatial CVE?

We also need to pose the question "are avatars always the best way to represent a user?". In Chapter 11 it is argued that there are cases in which avatars simply get in the way and a simpler representation is more effective. The authors also argue for the use of additional media such as video running in parallel with the CVE to give a richer understanding of the actions of remote participants. This argument is further reinforced in Chapter 2, in which the authors argue that it is necessary to provide a stronger integration between CVEs and the use of other media.

1.3.2 I Know What I See, but What Do *You* See?

In the real world people have very sophisticated spatial reasoning skills. These are used to allow people to perform tasks effectively but also to enable social situations to occur smoothly with the minimum of explicit (verbal) negotiation. In comparison with the real world, virtual environments are impoverished in terms of the information that they feed to our senses – a person in a virtual environment is reduced to "tunnel vision" and sees a much simplified world without many of the cues that are (unconsciously) so helpful in the real world, such as shadows, subtle sounds, rich lighting and colours, and proprioception. This sensual lack can lead to a number of problems for the virtual inhabitant that may not be obvious at first glance. Chapters 5 and 7 highlight some of these problems and the ways in which people attempt to cope with them. In this section we have two chapters that consider these issues from a more technical viewpoint. Chapter 9 describes Kansas, a flat 2D world in which the presentation is so ordered that it is easy for people to deduce what is seen by others – in fact the views of others are superimposed on the world in a 2D equivalent of the visible view frustums proposed in Chapter 7. Although in Kansas some users can see objects not available to others, when users' view areas overlap it is clear who can see what and that all users see the objects in the same way.

However, given that virtual environments are not constrained by the real world there are cases in which it might use useful to take advantage of the fact that users don't necessarily see the same things – something which is usually explicitly minimized in CVE systems. Chapter 8 demonstrates how subjective elements can be explicitly added to CVE systems in order to let people customize their view while still remaining in contact with other people – who, except in the case of people suffering from sensory dysfunction, does not occur in the real world. This theme of subjectivity is also taken up in Chapter 12, which also illustrates that it is not always desirable for every user to experience the world in the same way.

1.3.3 Technical Issues and System Challenges

Unfortunately we must live in a (real) world with imperfect hardware and physical constraints. The speed of light imposes a minimum time to send a message to someone – while we normally do not notice such delays, if the person is on the other side of the world these delays can become perceptible. Even worse, we are not just limited by the speed of light but by network performance, which means that we will often experience delays far worse than simply those imposed by the speed of light. And this is not all: the real world is of such richness that a computer does not yet exist with the performance to even approach a realistic modelling and rendering of it. Finally, despite the hype surrounding VR current interface devices are often anything but *natural* to use; our day-to-day skills from interacting in the physical ("real") world do not easily transfer to interactions with and within VR worlds. With all these limitations where does that leave us? How can we make the best of what we've got? Given that we can't have perfection, can we at least choose where to have the imperfections? It is these and other similar issues that the chapters in this section will try to respond to. Chapter 3 provides a road map of the way ahead for CVE systems and indicates where designers need to focus their attention in future. Chapter 4 shows us how, in a world of imperfect networking, we can choose where to make our trade-offs.

Historically, VR systems were conceived as being worlds apart, not only in what they presented to the user but also in the ways they interacted with the rest of the user's computing systems. If you wanted to discuss a document in a meeting taking place in a CVE then that document had to have an application that could represent it in the CVE. This is one factor implicated in the slow uptake of CVE systems in general use, and Chapter 2 shows why this way of organizing a CVE system is such a mistake.

1.3.4 So, Now We're in a CVE What Do We Do?

So far, the chapters in this book have presented a number of features of CVEs, both negative and positive, and shown some of the impact that the underlying system

Figure 1.3 Paris from "Le Deuxieme Monde" (the Second World) views of virtual Paris. (Copyright Canal+, All Rights Reserved.)

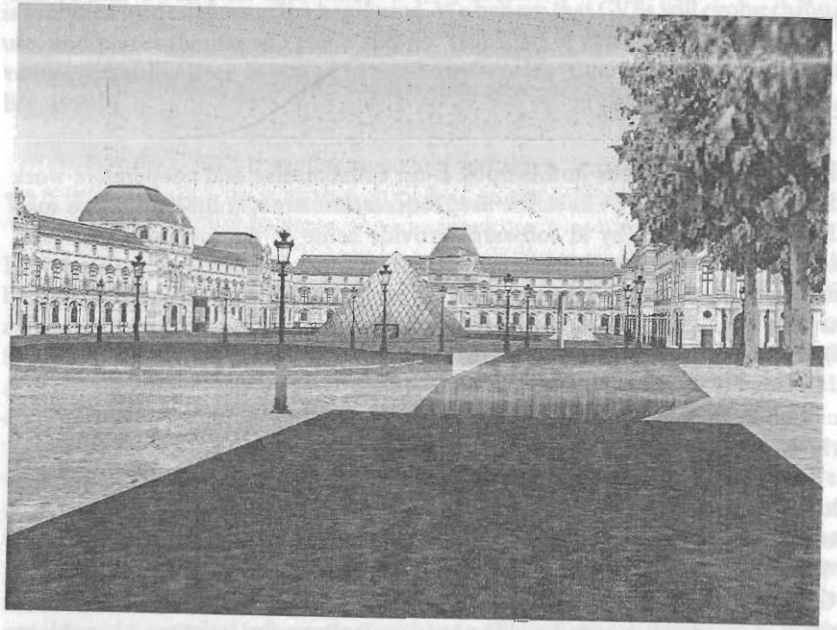

Figure 1.4 Paris again, from "Le Deuxieme Monde" (the Second World) views of virtual Paris. (Copyright Canal+, All Rights Reserved.)

22 000 "citizens" who meet daily, participate in discussions and benefit from events staged in the environment. Users have a free choice of the appearance of their avatars. More than being simply visitors, the citizens can build their own personal spaces while CANAL+ takes care of the representation of the principal districts of Paris. Celebrities are regularly invited to take part in events in the Second World, such as games, quizzes and discussion forums. There are also clubs organized by the inhabitants of the Second World and exhibitions organized by other organizations, such as the galleries of the Louvre in 3D created by employees of the museum.

The Second World is interesting in another respect because the community created is not simply virtual. Second World citizens regularly stage meetings in the real world and can travel large distances to take part. Such crossover is often seen in communities that build up over time and where interactions are ongoing (e.g. Mynatt *et al.*, 1998, 1999; Preece, 1999).

1.3 Themes Covered in the Book

1.3.1 Bodies, Presence and Interactions

In a spatial CVE it is almost a given that a user is represented by an avatar; however, there are many problems which may not be obvious at first, such as how well users can understand someone's actions and point of view judging solely by their view of that person's avatar. Chapters 6 and 7 highlight some of the problems with avatars in collaborative spaces and show two differing ways in which to control one's avatar. In Chapter 7, the user deals at a fine-grained level with every aspect of his or her avatar, and we can see both the opportunities and difficulties inherent in this approach. Granted, one has fine-grained controllability, but in the circumstances, owing to a number of factors such as lack of view, this is not as powerful as it might be.

We can see from Chapter 3 that it is almost an impossibility to get true real-time pictures of one's own and another's avatar. If it were possible, we would in effect be breaking the physical limits of the speed of light and consequently electrons travelling in a network. Any solution to this in terms of software fixes, such as anticipating user behaviour, will have, as Chapter 3 points out, implications for the granularity of our control of the avatar, whether physical (e.g. a robot in the real world) or virtual. Chapter 3 points out the very real implications for other fields, such as tele-robotics.

Chapter 6 addresses this question from a different angle, asking, in effect, why we would want to control our avatar closely in the first place. The authors liken the avatar to a motor vehicle, and note that when we drive, our cognitive processes are largely taken up with driving. It is often difficult, they say, even to hold a conversation with another person in a car if one is driving at the same time. Looking at it this way, they wonder how we are meant to "drive" avatars *and* also do a number of different things in the space, such as the banal everyday tasks of holding conversations or collaborating on a piece of work. The solution is, in their words, to "let the

activities, which in turn provides a sense of shared context. In many CVEs, it is also possible to leave text, audio or video messages that can be played back.

Awareness of Others

There are several ways of conceptualizing "awareness". For example, Dourish and Bellotti (1992) state that awareness is an "understanding of the activities of others, which provides a context for your own activity". Such a view of awareness centralizes intentional awareness. Tacit or background awareness involves consideration of peripheral as well as focused attention and more accurately characterizes what occurs when team members are engaged in parallel but independent ongoing activities. In contrast to intentional awareness, such tasks often require moment-to-moment peripheral coordination (Kendon, 1990). For example air traffic controllers are not aware of each other's moment-to-moment activities, but peripheral vision and background sound all provide information such that if a disruption occurs, unplanned collaborative activities can ensue (see also Heath and Luff's (1996) analysis of London Underground workers). Some consideration of providing such moment-to-moment awareness or sense of co-presence exists in work on video portholes and video tunnels where offices are linked with video cameras and monitors (Buxton and Moran, 1990; Dourish *et al.*, 1996; Gaver, 1992). Awareness can also relate to activities outside of the current task context where one is interested in the activities of a collaborator who is not currently present and who may not be working on the shared task. Often we need to know where to get hold of someone and/or need to adjust our plans on the basis of when someone will be back. In everyday life, voice mail, answering services, answer phones and vacation email messages play this role.

Negotiation and Communication

Conversations are crucial for negotiation and communication about collaborative activities. Collaborative work requires the negotiation not only of task-related content, but also of task structure in terms of roles and activities and task/sub-task allocations. Further, informal conversations underpin the social fabric that sustains ongoing collaborative relationships. Studies of conversation have demonstrated the extent to which explicit hand gestures and head nods, eye gaze and eyebrow raises determine the way in which utterances are interpreted[1]. As well as these "backchannel" gestures that punctuate and give meaning to many conversational utterances, cues from dress, posture and (often culturally determined) mannerisms provide much of the background context for our interpretations of verbal negotiations (Hewstone *et al.*, 1988; Kendon, 1990).

In graphical CVEs such backchannel gestures are often hard to achieve with embodiments where nuanced subtle gesturing is not easily supported. This can

[1] See Cassell *et al.* (2000) for a number of chapters that detail nonverbal channels of communication.

cause disruptions in conversational flow (see Chapter 7). Further, until it is easier for the users of CVEs themselves to fashion embodiments easily, embodiments will not be tailorable and will not offer as many options for giving personalized indications of one's cultural and personal identifications. This has been raised as an issue in the context of representations of gender and race within the broader cyberculture in which CVEs sit (Kolko *et al.*, 2000). In text-based CVEs such cues are often provided explicitly in text (see Chapter 13 for a consideration of cultures in text-based CVEs).

Flexible and Multiple Viewpoints

Tasks often require use of multiple representations, each tailored to a different points of view and different subtasks. For example, Bellotti and Rogers (1997) offer a detailed analysis of the production of a daily newspaper. In this process many different representations are used to design the layout of the paper; these vary from hand-drawn to computer-generated and reflect different task requirements. In certain cases, one individual may require multiple representations to reflect different aspects of their task(s), while in other cases different individuals may require tailored representations to provide information specific to their tasks. In many of the chapters in this volume authors detail provision of multiple views and the support of alternative representations to suit the different roles people are taking and the different information they require (e.g. Chapters 8 and 9; and Greenhalgh and Benford, 1995).

1.2.3 Recreational CVEs

Although our main focus in this book is on collaborative and cooperative work within CVEs, of course CVEs can be used for recreation too. Multi-user games such as Doom and Quake (by id software) provide some of the most effective (and violent) demonstrations of the use of commodity hardware to produce an immersive experience. Watch someone playing Quake and you will see them move their body in an attempt to dodge fire from their opponents even though this (physical) action has no effect on their virtual body. Tap such a person on the shoulder and they will typically almost leap out of their seat – such is the shock of the forced transition from being absorbed in the virtual world to being dragged into the real world. Not all recreational CVEs need be violent though – Chapter 10 describes how "the blob", an amorphous pulsating virtual object, can provide a way to help young children create stories.

Another example of a recreational but non-violent CVE is the "Second World" (Le Deuxième Monde), the first European 3D virtual community, which is hosted on CANAL+'s Web site (`http://virtuel.cplus.fr/`). The Second World takes the form of a virtual Paris rendered in photo-realistic quality (see Figs. 1.3 and 1.4). Access to the Second World is open to everyone, requiring a simple download of a Web browser plug-in from Blaxxun. The Second World already has more than

"home" notions above, places can be public or private for example (see de Certeau *et al.*, 1998, on concepts like the "neighbourhood"). Some technologies, such as 3D virtual environments, are inherently spatial, and it could be argued that some 3D virtual environments suffer because they depend too much on *space* without supporting the creation of a sense of *place*. CVEs need to support the *evolution of places for interaction*, and not simply provide spaces where interaction can take place, if they are to be successful in the long term.

Further, if we are to support multiple groups within a virtual space, appropriate landscapes in which varying activities can occur must be designed and developed. These digital landscapes will have many differing places for interaction, supporting many concurrent ongoing activities with multiple groupings of people. Sometimes the places will be bordered by digital walls or separated through the use of distance (see Chapter 15 for a consideration of virtual walls and boundaried digital spaces). That is not to say such places can be reliably *designed* or predetermined. As noted by Roderick Nash (1973; cited in Crandell, 1993) when discussing the evolution of the physical landscape:

> No group sets to create a landscape, of course. What it sets out to do is to create a community, and the landscape as its visible manifestation is simply the by-product of people working and living, sometimes coming together, sometimes staying apart, but always recognizing their interdependence.

Similarly, Harrison and Dourish (1996) note the appropriation of technologies that occurred in many media space experiments. Often in such situations such appropriation means that, in the end, the use of the technologies differs from that which is *intended* by designers at the beginning. We believe that CVEs will evolve through use, and places therein will come and go. This kind of ebb and flow of activity in virtual places has been observed in text-based worlds (Cherny, 1999; Churchill and Bly, 1999b).

1.2.2 Collaborative Work in Virtual Spaces and Places

Elsewhere we have argued that CVE systems and applications must go beyond being "cool" and having aesthetic appeal, and will only become everyday places for action and interaction if they are designed to serve a purpose (Churchill and Snowdon, 1998). Specifically, in order to support collaborative and cooperative activities, it is important that virtual environments offer the means to access (task-) appropriate information as well as communication tools.

Although it is important not to try to simply replicate what we think of as "reality", when designing systems to support collaborative work we can learn a great deal from observations of people working and collaborating together in conventional settings (Bowers and Martin, 1999; Engeström and Middleton, 1996; Moran and Anderson, 1990; Heath and Luff, 1991). In all "real world" domains, collaborative work involves the interleaving of individual and group effort, so collaborative work involves considerable complex information exchange (Bellotti and Rogers, 1997; Harper, 1997; Heath and Luff, 1991; Heath and Luff, 1996; Hutchins, 1990; Hutchins and Klausen, 1996; Suchman, 1996). These interleaved, singular-to-shared activities

require considerable explicit and tacit communication between collaborators to be successful. Individuals need to negotiate shared understandings of task goals, of task decomposition and sub-task allocation, and of task/sub-task progress. It is important that collaborators know what is currently being done and what has been done in the context of the task goals. Until recently, most CVEs have been used as meeting places where group activities are the central task (e.g. Greenhalgh, 1997b). Recent work, however, has been aimed at supporting the kind of situation described above, where more complex, interleaved individual and collaborative activities can be carried out within one VE with team members moving continually between individual and collaborative activities (see Chapter 6).

We have noted that such observations point to a number of key features that software should aim to support (see Churchill and Snowdon, 1998). To summarize, these are:

- Shared context
- Awareness of others
- Negotiation and communication
- Flexible and multiple viewpoints

Sharing Context

Shared context is crucial for collaborative activities. "Shared context" can mean many things; it can mean shared knowledge of each other's current activities, shared knowledge of others' past activities, shared artefacts and shared environment. Together, these lead to shared understandings. Shared physical spaces and familiar places facilitate or "afford" shared understandings. As Gaver (1992) states

> in the everyday world, collaboration is situated within a shared, encompassing space, one which is rich with perceptual information about objects and events that can be explored and manipulated

Within these shared spaces, focused and unfocused collaboration is accomplished through alignment towards the focal area of the shared activity, such as a shared document (Heath and Luff, 1996), and through gestures like pointing toward portions of the document for added emphasis and clarity. When artefacts are shared, not only do they become the subject of communication between users, but also the medium of communication; as one user manipulates an object, changes to it are visible to others in an externalization of the processes of change. For the design of systems to support collaborative work, this means that shared artefacts should be visible and available for local negotiation (Dix, 1994; Heath and Luff, 1996) and often the current focus of attention should be indicated. This can drive subsequent activities. Where actions are not physically co-located and co-temporal, providing shared context is more difficult. It is also important to share context in asynchronous work collaborations through meeting capture, version control and so on; such tools provide activity audits and "awareness" of others'

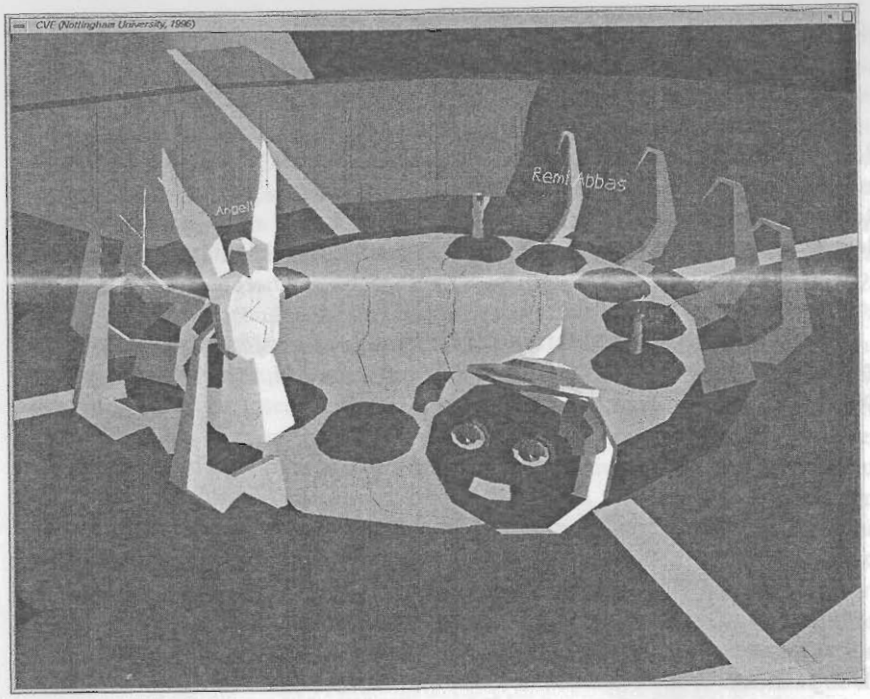

Figure 1.1 Embodiments in a digital landscape. From the event "MASSIVE" at the Nottingham NOWninety6 arts festival in London. (copyright Nottingham University, Nottingham UK, 1996. All Rights Reserved.)

an important distinction between space and place. A space is a physical (or digital) volume or container that can house artefacts and in which events can take place. However, one physical space can host several places over time; a sports hall can be a place for playing sport, a place for prize-giving, a place for a student play, a place for a jumble sale and a place for a dance class. Although the physical space is the same, there are multiple understandings of how that space is to be used, depending on the occasion. There are also multiple ways of interacting and behaving in that space depending on the occasion or activity for which the space is being used. In contrast to our notions of "space", "place" has inherent within it a notion of the activities that occur there – activities that *take place* there. The space only becomes a "place" when an understood activity is scheduled or ongoing. Thus, space and place do not share a 1:1 mapping (Munro, 1999). A space can also be seen as a different place at one and the same time, depending on how it is "read". For example, inner city gangs use and read graffiti as markers of territories for activity (Munro, 1999). These physical spaces will have entirely different meanings for the different gangs in terms of whether one is in "home" or "enemy" territory (e.g. Armstrong, 1988). The way the graffiti is read is partly to do with its location – if in the territory of the enemy gang, it is regarded as a great success by those committing the incursion

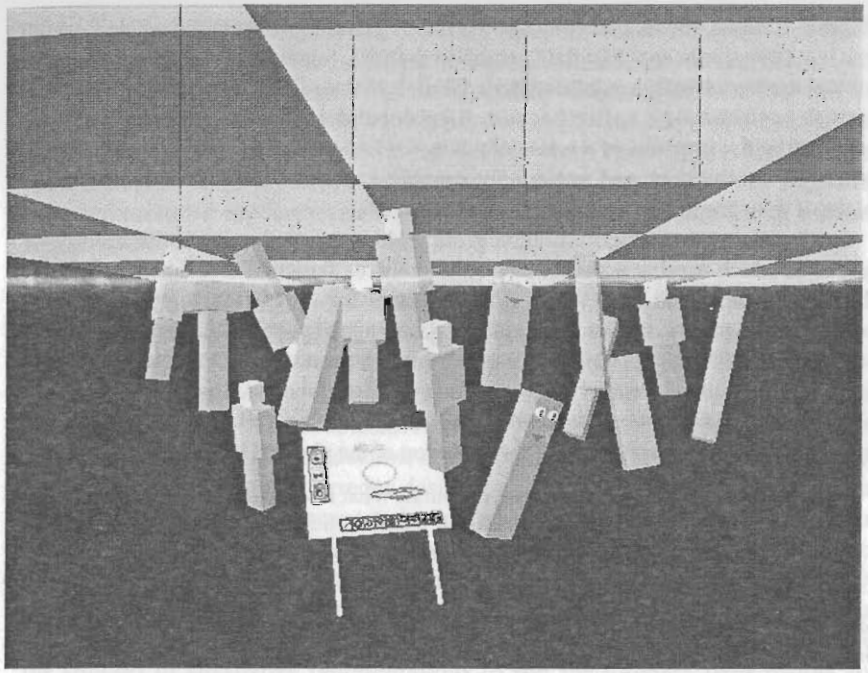

Figure 1.2 Virtual corpses below the digital "land" plane. (Copyright Swedish Institute of Computer Science, Sweden, 1996. All Rights Reserved.)

(Patrick, 1973). Thus a physical space may have many different meanings and modes of inhabitancy depending on one's reading of that space.

Some technologies, such as media-spaces, are focused on connecting different physical *places* without having a corresponding notion of spatiality, virtual or otherwise (Harrison and Dourish, 1996; see also Dourish *et al.*, 1996). Some cases, such as the CAVE system described in Chapter 12, illustrate how spaces and places can be combined in various ways – the CAVE is a real-world physical space that acts as a place in which users can experience virtual spaces and places. In the CAVE, users must negotiate both virtual and real spaces but can share their experiences with others in a way not possible in CVE systems based on head-mounted displays (HMDs), in which the user interface isolates users from the real world. Text-based CVEs can create spaces (rooms) that have only a textual (described) spatial geometry but which are certainly places in which people meet, talk, work and play (Churchill and Bly, 1999a,b,c and Chapters 13 and 14 of this volume). These virtual text-based rooms achieve their sense of place because they are used over time, because inhabitants have both planned and unplanned meetings there and because such meetings have a sense of continuity – the social conventions that make a space a place develop in these circumstances (see Chapter 14). Spaces become places when they are associated with living and lived experiences (Giddens, 1984). Places are imbued with meanings that underpin how we behave there – like the "enemy" and

research in these environments, the current state of the art, and where research and technical developments are heading. In addition, this book is multi-disciplinary – that is, rather than focus only on computer science issues, we have aimed to show how the field of CVE design and development is actively informed by other disciplines, such as psychology, sociology, work practice studies, architecture, artificial intelligence and art. We also intend to maintain a broad perspective in our choice of technology – there are many existing books on 3D Virtual Reality technology (e.g. Anders, 1998; Rheingold, 1992; Vince, 1998; Weishar, 1998) and we do not aim to duplicate them. Rather, we will show how other technologies, such as text-based virtual environments, can also be valid and usefully classified as CVEs (see Curtis, 1996; Curtis and Nichols, 1992; and Chapter 14 of this volume). Further, we would like to note that our focus is on the design and hands-on use of specific technologies. There are many fascinating books that theorize about the impact of these (and related) emerging technologies from social and socio-political perspectives (examples include Crang *et al.*, 1999; Porter, 1997; Kolko *et al.*, 2000). We point readers to these volumes for in-depth social and cultural analyses and critiques.

We have organized the book into a number of sections in order to group together related chapters. However, since CVEs are such an interdisciplinary field there are many links between chapters in different sections and the reader should bear in mind that this grouping is a guide and based on our personal judgement. In the remainder of this chapter we will introduce some of the background behind CVEs and the groupings that we have chosen.

1.1.1 What are Collaborative Virtual Environments? A Broad Definition

Information sharing is central to collaborative work. CVEs can help with information sharing and communication tasks because of the way in which they provide a context for communication and information sharing to take place. However, before we go deeper into the nature of CVEs it is first necessary to define what we mean by a *Collaborative Virtual Environment*. Here is one definition of a CVE that probably fits many people's expectations and is in line with the descriptions found in books such as *Neuromancer* (Gibson, 1989):

> Collaborative virtual environments are distributed virtual reality systems that offer graphically realized, potentially infinite, digital landscapes. Within these landscapes, individuals can share information through interaction with each other and through individual and collaborative interaction with data representation.

While this is not a bad definition, we have found it a little restrictive. Here is the definition that we have been using as the introduction to the current CVE conference:

> A CVE is a computer-based, distributed, virtual space or set of places. In such places, people can meet and interact with others, with agents or with virtual objects. CVEs might vary in their representational richness from 3D graphical spaces, 2.5D and 2D environments, to text-based environments. Access to CVEs is by no means limited to desktop devices, but might well include mobile or wearable devices, public kiosks, etc.

This definition makes it clear that although CVEs are normally associated with 3D graphical environments this need not necessarily always be the case. In the next sections we reflect further on what we mean by Collaborative Virtual Environments, focusing on the nature of environment and what it means for activity, and the nature of collaboration and collaborative work. Following this, we will outline the sections of the book.

1.2 Collaborative Virtual Environments: Some Considerations

1.2.1 Virtual Environments as Places for Action and Interaction

CVEs represent a shift in interacting with computers in that they provide a space that contains or encompasses data representations *and* users. Elsewhere we have noted (Churchill and Snowdon, 1998):

> CVEs represent the computer as a malleable *space*, a space *in which* to build and utilize shared places for work and leisure. CVEs provide a terrain or digital landscape that can be 'inhabited' or 'populated' by individuals and data, encouraging a sense of shared space or place. Users, in the form of embodiments or avatars, are free to navigate *through* the space, encountering each other, artefacts and data objects and are free to communicate with each using verbal and non-verbal communication through visual and auditory channels.

Figure 1.1 shows a virtual landscape that consists of a central circle or platform. Two avatars or embodiments are currently inhabiting the virtual space – an angel and a round, floating, face. Both of the embodiments can navigate through the space; they can see each other and can both see the landscape clearly. They can interact through audio connections and can orient their bodies to each other. Evidence suggests that people often orient their virtual embodiments toward each other in ways that approximate real-life interactions (Smith *et al.*, 2000). A number of chapters in this volume detail issues involved in using such embodiments within virtual space to interact with others and with virtual objects (for example see Chapters 5 and 7).

Of course, being a digital landscape, it is possible for digital embodiments to behave in ways that our physical counterparts cannot: the rules of "real world" physics do not apply here. In Fig. 1.2 we see a number of virtual embodiments that have floated seamlessly through the floor and now reside *below* the ground plane. Much work is currently ongoing considering what the correct models are for designing collision detection between virtual objects including embodiments, walls and floors within virtual worlds. In this volume, Chapter 3 offers a discussion some of the issues involved in creating such virtual world physics.

However, it takes more than provision of the digital landscape to create digital spaces places that are fit for human–human interaction. Indeed, there has been much debate in recent years about the very nature of space and place in physical and in digital settings (Harrison and Dourish, 1996). Although virtual environments are often assumed to be spaces and most utilize spatial metaphors, there is

Chapter 1

Collaborative Virtual Environments: Digital Spaces and Places for CSCW: An Introduction

Dave Snowdon, Elizabeth F. Churchill and Alan J. Munro

1.1 Introduction

In the late 1980s Virtual Reality (VR) burst onto the public stage propelled by a wave of media interest and related science fiction novels such as *Neuromancer* by William Gibson (Gibson, 1989). VR promised to revolutionize the way in which we experience and interact with computers, and research into the field mushroomed. More recently, the hype surrounding VR has died down and, although it is receiving less public attention, serious work is continuing with the aim of producing useful and usable technology. At the centre of current work related to VR is the field of Collaborative Virtual Environments (CVEs). This field has as its goal the provision of new, more effective means of using computers as tools for communication and information sharing with others. Many CVE systems have been constructed. Some of these are desktop systems and applications; but large public virtual spaces have also been constructed (such as AlphaWorld at `http://www.activeworlds.com/`; see Chapter 15). CVEs are also being used to experiment with new forms of art and interactive television (Benford *et al.*, 1997a,b; Benford *et al.*, 2000a).

In 1996 the first international conference on CVEs took place in Nottingham, UK – this received an enthusiastic reception and was followed in 1998 by CVE'98. Two special issues of the journal *Virtual Reality. Research, Development and Applications* on Collaborative Virtual Environments followed, published in 1998 and 1999. At the time of writing, CVE 2000 is being planned and will take place in San Francisco in September 2000, with sponsorship from the Association of Computing Machinery (ACM – `http://www.acm.org/`). Given this increasing interest, we felt that now is an appropriate time to collect together a book that can serve as an introduction to this broad and rapidly developing field.

In editing this book, we aim to give a broad introduction to Collaborative Virtual Environments and how they can bring people together and allow them to communicate. By communication we do not just mean text or audio conversations, but we also wish to consider how artefacts and embodiments are also an essential aspect of communication, and thus how the representation of people and artefacts such as documents and tools within CVEs can facilitate communication. Our goal is for readers to get a better feel for why CVEs are an interesting subject, where CVEs can be applied to support collaborative applications, the important aspects of the