Awareware: Narrowcasting Attributes for Selective Attention, Privacy, and Multipresence

Michael Cohen and Owen Noel Newton Fernando

1 Introduction

The domain of CSCW, **c**omputer-**s**upported **c**ollaborative **w**ork, and DSC, **d**istributed **s**ynchronous **c**ollaboration, spans realtime interactive multiuser systems, shared information spaces, and applications for teleëxistence and artificial reality, including collaborative virtual environments (CVEs) [9]. As presence awareness systems emerge, it is important to develop appropriate interfaces and architectures for managing multimodal multiuser systems. Especially in consideration of the persistent connectivity enabled by affordable networked communication, shared distributed environments require generalized control of media streams, techniques to control source → sink transmissions in synchronous groupware, including teleconferences and chatspaces, online role-playing games, and virtual concerts.

There are two main techniques currently used for managing information in contemporary systems to address the problem of human information overload—proximity-based filtering, as used by many games and as formalized by the Benford et al. model described below, and explicit **d**egree-**o**f-**i**nterest (DoI) filtering, as seen in buddy lists of **i**nstant **m**essaging (IM) systems or clan-based chat channels of **m**assively **m**ultiplayer **o**nline **r**ole-**p**laying **g**ames (MMORPGs).

Anticipating ubicomp networked appliances and information spaces, we are exploring the integration of various multimodal (auditory, visual, haptic) I/O devices into mixed and virtual reality groupware suites. Such environments are characterized, in contrast to general multimedia systems, by the explicit notion of the position (location and orientation) of the perspective presented to respective users; often such vantage points are modeled by the standpoints and directions of icons in a virtual

Michael Cohen
Spatial Media Group, University of Aizu, e-mail: mcohen@u-aizu.ac.jp

Owen Noel Newton Fernando
Mixed Reality Lab., National University of Singapore, e-mail: newtonfernando@mixedrealitylab.org

space. These icons might be more or less symbolic (abstract) or figurative (literal), but as representatives of human users, are therefore “avatars” [6]. Avatars reify embodied virtuality, treating abstract presence as a user interface object.

This chapter reviews the basic and extended notions of awareness and presence in virtual environments, explains the idea of multipresence, surveys related models of groupware awareness, and presents a formalization of narrowcasting, which ideas are deployed in two integrated interfaces, for workstations and mobile phones, considered as case studies.

1.1 Presence, Telepresence, and Copresence

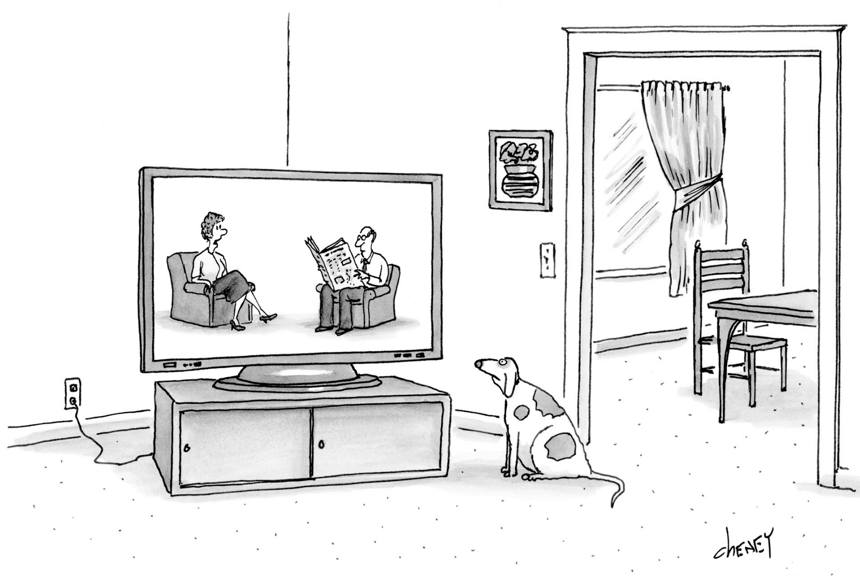

“Notice anything different?”

Fig. 1 Metapresence. (© The New Yorker Collection 2006 Tom Cheney from `cartoonbank.com`. All rights reserved.)

Presence may be the most elementary component of virtual collaboration. It has been used broadly, but generally presence is the feeling of “being there,” as disquietingly suggested by Figure 1. Slater et al. [51, 52] define presence as “a state of consciousness, the (psychological) sense of being in the virtual environment.” In the context of environment, presence describes the degree to which one feels a part of some virtual space— that the space exists and one is occupying it.

The term "telepresence" has been used in industry since scientists and engineers started to design and develop remote control systems and industrial robots. Generally, telepresence [8] allows users to experience a physical space through display and control interface elements connected to remote sensors and actuators.

"Copresence" [50] is a sociological concept that describes the way people interact with each other. Copresence is primarily used to refer to either telepresence (the sense of being together with other people in a remote physical environment) or virtual presence (the sense of being together with other people in a technology-generated environment).

The sense of copresence is different from the mode of copresence: whereas mode of copresence refers to one's spatiotemporal collocation with others, sense of copresence involves one's perceptions and feelings of being with others. One's sense of being with others is basically a psychological phenomenon, which may or may not correspond to the actual state of copresence. An individual, for example, can be made to feel that he or she is interacting with another human, even through the individual is in fact completely alone. Psychological states— such as mood, alertness, and prior experiences— affect one's sense of copresence, and environmental factors— such as temperature, light, sound, and smell— may also influence one's sense of being with others.

1.2 Awareness and Presence Awareness

The concept of 'awareness' has been used in numerous ways in the **h**uman-**c**omputer **i**nteraction (HCI) [40] and CSCW literature [29]. Awareness has been defined [24] as "an understanding of the activities of others, which provides a context for [one's] own activity." This definition encompasses many kinds of displays of colleagues' actions in shared information spaces— such as whiteboards, alerts about people's status, access privileges to information, prior actions, and so on.

Presence awareness provides information like the location, identity, activities, and neighbors of someone. A wide range of distributed applications requires presence awareness [12] [58, 23], including instant messaging (IM) systems, and groupware applications like chat, audio- or video-conferencing systems. Well-known current applications include messengers like those from AOL, Yahoo!, or MSN, as well as CSCW applications and virtual 3D communities like Active Worlds or Second Life. Currently, presence awareness is mostly used for IM systems to let users know when others, especially those on contact (or buddy) lists, are on-line and willing to accept messages. Presence-aware groupware applications are sensitive to the receptiveness of the respective participants. When a messaging system is part of an integrated communications platform, presence awareness can become more sophisticated. It can notify others when a user is on-line, willing to accept phone calls at home, or has a mobile phone turned on [38]. A conferencing system might know that a session member is asleep, and not awaken her for non-urgent realtime voice chat. Users might even set presence messages so others trying to contact them will learn that they've gone out for while and will return at a certain time.

1.3 Narrowcasting and Privacy

In analogy to uni-, multi-, and broad-casting, "narrowcasting' refers to the deliberate filtering of multiple duplex information streams, a model for rich CSCW and social networking. Narrowcasting systems extend broad- and multicasting groupware systems by allowing various information streams to be filtered, for privacy, security, and user interface optimization. The narrowcasting operations described in this chapter suggest an elegant infrastructure for such collaborative environments, an idiom and service for selective attention and presence awareness. For simple example, a user's voice might by default be shared with all others in a chatspace, but an appropriate interface would allow a secret to be shared with some select subset ("inner circle") of the session members.

Traditional mixing idioms for enabling and disabling various audio sources employ `mute` and `solo` functions which selectively disable or focus on respective source channels. As summarized in Table 1, which previews the case studies presented later in this chapter, sinks are defined as duals of sources in virtual spaces [15], logical media stream receivers. Exocentric interfaces, which explicitly model not only sources, but also position and multiplicity of sinks, motivate the generalization of audio mixer commands `mute` & `select` (or `cue` or `solo`) to `exclude` and `include`, manifested for sinks as `deafen` & `attend`, a narrowing of stimuli by explicitly or implicity blocking out and/or concentrating on selected entities, as elaborated by Figure 2.

The general expression of activation is

$$\texttt{active}(x) = \neg\texttt{exclude}(x) \wedge (\exists\, y\, (\texttt{include}(y) \wedge (\texttt{self}(y) \Leftrightarrow \texttt{self}(x))) \Rightarrow \texttt{include}(x)). \quad (1)$$

So, for `mute` and `select`, the relation is

$$\texttt{active}(\text{source } x) = \neg\texttt{mute}(x) \wedge (\exists\, y\, (\texttt{select}(y) \wedge (\texttt{self}(y) \Leftrightarrow \texttt{self}(x))) \Rightarrow \texttt{select}(x)), \quad (2a)$$

`mute` explicitly turning off a source, and `select` disabling the collocated (same window/room/space) complement of the selection (in the spirit of "anything not mandatory is forbidden"). For `deafen` and `attend`, the relation is

$$\texttt{active}(\text{sink } x) = \neg\texttt{deafen}(x) \wedge (\exists\, y\, (\texttt{attend}(y) \wedge (\texttt{self}(y) \Leftrightarrow \texttt{self}(x))) \Rightarrow \texttt{attend}(x)). \quad (2b)$$

Fig. 2 Formalization of narrowcasting and selection functions in predicate calculus notation, where '¬' means "not," '∧' means conjunction (logical "and"), '∃' means "there exists," '⇔' means mutual implication (equivalence), and '⇒' means "implies."

Narrowcasting functions, which filter stimuli by explicitly blocking out and/or concentrating on selected entities, can be applied not only to other users' sources

Table 1 Roles of $^{s}\mathrm{OU}^{\mathrm{rce}}_{\mathrm{Tput}}$ and $^{s}\mathrm{IN}^{\mathrm{k}}_{\mathrm{put}}$: Iconic and figurative attributes of narrowcasting functions extend avatars to denote invoked filters.

	Source	Sink
Function	radiation/transmission	reception
Level	amplification	sensitivity
Direction	OUTput	INput
Instance	speaker	listener
Transducer	loudspeaker	microphone or dummy-head
Organ	mouth	ear
Tool	megaphone, bullhorn	ear trumpet
Exclude	`mute`	`deafen`
Inhibit in ι·Con	$\overset{-}{\triangle}$	$-\triangle-$
Inhibit Self in Multiplicity *reflexive*	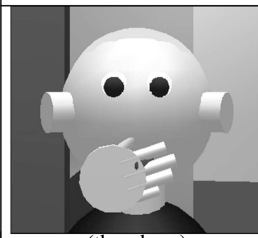 (thumb up)	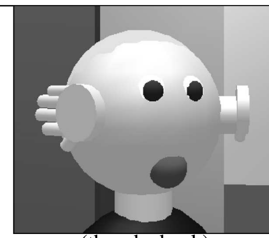 (thumbs back)
Inhibit Other in Multiplicity *transitive*	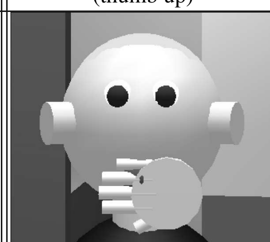 (thumb down)	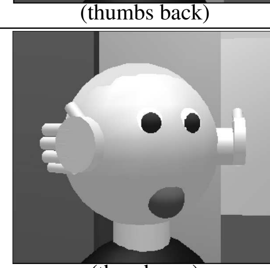 (thumbs up)
Include	`select` (`solo` or `cue`)	`attend`: `confide` and `harken`
Assert Target in ι·Con	$\overset{+}{\triangle}$	$+\triangle+$
Assert Target in Multiplicity *explicit*	 (megaphone)	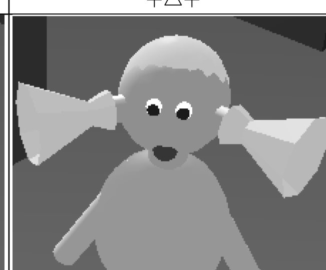 (ear trumpets)
Assertion Side-effect in Multiplicity *implicit*	 (translucent hand)	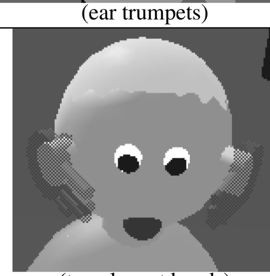 (translucent hands)

Fig. 3 The price of privacy. (© The New Yorker Collection 1996 Sam Gross from cartoonbank.com.)

and sinks for privacy, but also to one's own, for selective attendance or presence. "Privacy" has two interpretations, as suggested by Figure 3. The first association is that of avoiding "leaks" of confidential information, protecting secrets. But a second interpretation means "freedom from disturbance," in the sense of not being bothered by irrelevance or interruption. The distributed interfaces described in this chapter feature narrowcasting operations that manage privacy [1] in both senses, by filtering duplex information flow.

The inclusion and exclusion narrowcasting commands for sources and sinks are like analogs of burning and dodging (shading) in photographic processing. The duality between source and sink operations is tight, and the semantics are identical: an object is inclusively enabled by default unless, a) it explicitly excluded (with `mute` [as in Figure 4] or `deafen`), or, b) peers of the same `self`/non-`self` class are explicitly included (with `select` or `attend`) when the respective object is not.

That is, if any avatar has been `select`ed, non-`select`ed avatars are implicitly `mute`d if their `self`-designation state is the same as a `select`ed avatar. In the same manner, if any `attend`ed avatars are in a given space, non-`attend`ed avatars are implicitly `deafen`ed if they are in the same class (`self` or non-`self`) as an `attend`ed avatar.

These narrowcasting attributes are not mutually exclusive, and their dimensions are orthogonal. Because a source or sink is active by default, invoking `exclude`

"With your kind permission, I've taken the liberty of putting Marvin on 'mute.'"

Fig. 4 Social narrowcasting. (© The New Yorker Collection 2008 J.B. Handelsman from `cartoonbank.com`. All rights reserved.)

and `include` operations simultaneously on an object results in its being disabled. For instance, a sink might be first `attend`ed, perhaps as a member of some non-singleton subset of a space's sinks, then later `deafen`ed, so that both attributes are simultaneously applied. As audibility is assumed to be a revocable privilege, such a seemingly conflicted attribute state disables the considered sink, whose audition would be restored upon resetting its `deafen` flag. Symmetrically, a source might be `select`ed and then `mute`d, akin to making a "short list" but relegated to backup.

1.4 Multiple Spaces

Non-immersive perspectives in virtual environments enable flexible paradigms of perception, especially in the context of frames-of-reference for conferencing and musical audition. In the model described in this chapter, users designate one or more avatars as "self," to establish correspondence between human pilots and virtual presence. Such exocentric interfaces, which explicitly include a representation of the subject as a "full citizen," allow users to perceive themselves "out of body," as well as the juxtaposition of multiple spaces.

Dix et al. [22] described their experiences of spaces in the Equator project,[1] in particular the way in which multiple spaces, both virtual and physical, can co-exist. People and objects may have locations in and relationships to both physical space and one or more virtual spaces, and these different spaces together interact to constitute overall system behavior and user experience.

The narrowcasting model mediates interactions between virtual and physical spaces to allow users to have presence in multiple places simultaneously. The workstation- and mobile-based interfaces described later in this chapter both support multiple spaces to enhance multipresence-enabled conferencing capabilities in CVEs.

1.5 Multipresence, Anycasting, and Autofocus

A humans is indivisible, and a person cannot physically be in multiple places at once. However, a unique feature of the interfaces described here is the explicit ability of a user to delegate several representatives simultaneously, increasing *quantity* of presence [14]. Such multipresence enables us to overcome some fundamental constraints of this human condition. Our interfaces encourage multipresence [12], by supporting `self`-designated avatars in multiple places simultaneously— allowing, for example, a user to monitor several spaces at once, refining the granularity of control.

Multiple sources are useful in directing one's remarks to specific groups. Multiple sinks are useful for monitoring several places at one, especially in situations in which a common environment implies social inhibitions to rearranging shared sources like musical voices or conferees, as well as individual sessions in which spatial configuration of sources, like the arrangement of a concert orchestra, has mnemonic value.

An "anycast" is a transmission between a sender and one of possibly several receivers on a network. The term exists in contradistinction to "multicast," transmission between a sender and multiple receivers, and "unicast," transmission between a sender and a single receiver. An anycasting service uses some criteria to choose a "best" or single destination from a set of candidates. We apply the same idea, finding the best sink (the one for which the source is loudest) for each source using an "autofocus" technique, illustrated by Figure 5 and described following.

In an audio interface that composites soundscapes from sinks' perspectives, the apparent paradoxes of multipresence, having avatars in multiple places or spaces

[1] `www.equator.ac.uk`

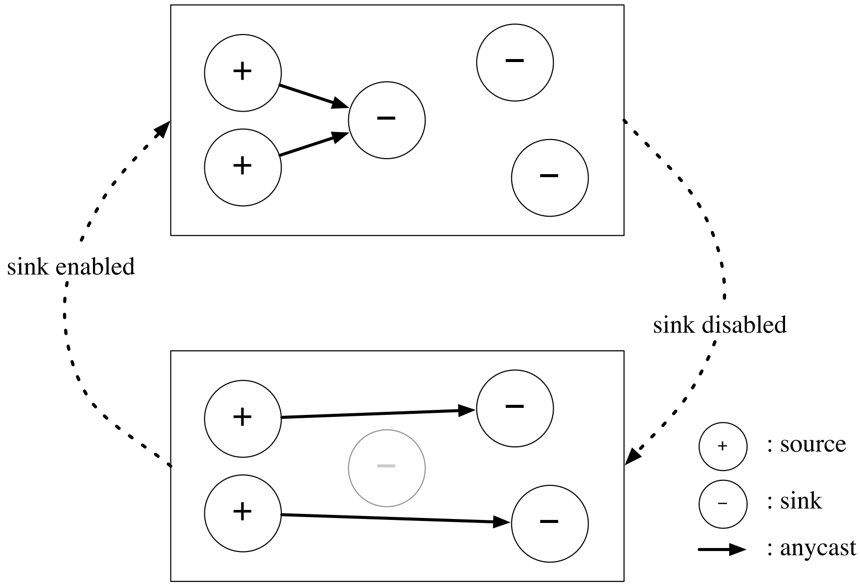

Fig. 5 Autofocused source → sink transmissions: If an intercepting auto-focused sink is `deafen`ed (or peers `attend`ed), remaining sinks adopt orphaned, anycasting, sources.

simultaneously, are resolvable by such an autofocus feature, which uses reciprocity, logical exchangability of source and sink, to project overlaid soundscapes and simulate the precedence effect to consolidate the audio display. If the sinks are distributed across separate conference rooms, each source is localized only with respect to the colocated sink. If multiple sinks share a single space, an autofocus algorithm is employed by anticipating "the rule of the first wavefront" [61, 32, 10, 27], the tendency to perceive multiple identical simultaneous sources from different locations as a single fused source. Rather than adding and averaging the contribution of each source [21, 55] to possibly multiple sinks, our system localizes each source only with respect to its respective best sink, consolidating the display. Multipresence encourages the narrowcasting-articulated audition (for sinks) or address (for sources) of multiple others.

1.6 Multipresence Scenarios

Most user interfaces support clipboard operations, `cut`/`copy`/`paste`. Such clipboard operations, enabled by dynamic deletion and creation of avatars, can be used for teleporting (`cut` & `paste`) and cloning (`copy` & `paste`) in CVEs, allowing avatars to convey narrowcasting narrowcasting attributes across multiple spaces in distributed applications and heterogeneous sessions.

A simple teleconferencing configuration typically consists of several icons, representing distributed users, moving around a shared conference space. Multipresence systems allow users to fork themselves, effectively increasing their attendance in virtual environments. For example, in a chatspace, a user might choose to designate two avatars as `self`, one to stand near an avatar corresponding to the user's mate, and another, perhaps on the other side of the 'room,' to stand near an avatar corresponding to the user's friend. Each of these avatars enjoys a local perspective, a situation awareness encompassing where the respective conversationalists are relative to the (no longer unique) `self`-associated avatar, as manifested visually and auditorilly.

In a parallel virtual rock concert, for continued example, a listener might want to pay close attention to both the drum and rhythm guitar, avoiding rearranging the instruments around a singleton sink to maintain consistency with other attendees. An active listener could fork her presence as `self`-designated avatars, locating one avatar near the drum, and `past`ing another near the guitar. To focus on only the previously-described chatspace, the listener might `attend` her avatars in it, so her other sinks in the virtual concert would be implicitly `deafen`ed. If she wanted to direct her voice privately to friends in either space, she might `attend` the friends' avatars, thereby implicitly `deafen`ing the others'.

2 Related Research

Gutwin et al. [31, 30] and Benford et al. [7] provide contrastable models of awareness issues in groupware. Gutwin et al.'s model attempts to identify the constituent information sources that communicate awareness and provide a framework that can assist the evaluation and design of awareness capabilities in groupware [13]. Benford et al. derive a model for awareness and interaction in virtual environment which focuses on the information space in mediating awareness. Each of these models is briefly described below. Benford et al.'s model (which is also considered in the "Abstractions of Awareness" chapter [39]) is quite similar to our narrowcasting model, so we review and compare it with the narrowcasting idioms more completely in section § 4.

Gutwin et al.'s model is derived from a top-down decomposition of awareness types, with particular emphasis on "workspace awareness," defined as awareness of others that is mediated by, or closely related to, actions on or around a shared workspace. This model identifies four types of awareness, which are specifically applied to group work dynamics:

Informal Awareness regards the sense of community among a group of colleagues [59, 60].

Conversational Awareness regards backchannels of communication that contextualise interaction.

Structural Awareness regards the protocols and structures used to formalize collaboration.

Workspace Awareness regards the capabilities of the media of collaboration, “the awareness mediated by the workspace.”

Benford et al.’s “spatial model of interaction” describes interaction based on metaphorically physical properties of space. The ability of a subject to perceive an object is affected by distance, direction, and possible obstructions. The key awareness abstractions in this model are “aura,” “focus,” and “nimbus”:

Aura is the portion of space in which interaction is enabled and allowed.
Focus (“attention function”) is the cummulative scope of regard. The more an object is within one’s focus, the more aware one is of it.
Nimbus (“publicity function”) is an object’s projection, its extent of exposure. The more a subject is within one’s nimbus, the more aware it is of one.

“Massive” [28] is a distributed virtual reality system providing facilities to support user interaction and cooperation via text, audio and graphics media, and interaction is controlled by these spatial models of interaction. The particular emphasis of Massive was on large-scale multiuser virtual environments, i.e. environments which might eventually support hundreds or thousands of simultaneous users. Aims of the Massive project (and its spatial model) were to provide rich forms of interaction which draw upon real-world behavior to make them useful and controllable in highly populated virtual worlds.

A third relevant model considers the form of the computing platform. With the spread of wireless communication and the desire to ‘travel light,’ collaboration across PCs and mobile devices (PDA, mobile phones, etc.) [48] [47] [41] [45] is a likely trend for future groupware applications. Velez et al. [57] investigated performances and communication patterns when collaborators use unequal computer platforms for their collaboration. They explored whether people use the same type of platform (homogeneous) or different platforms (heterogeneous) for communication, mainly considering PCs (**p**ersonal **c**omputers) vs. PDAs (**p**ersonal **d**ata **a**ssistants) for heterogeneous platforms. Their findings suggest that limited device capabilities can affect who is actually in charge and attention must be paid to the types of representations used on the mobile platform, as poor representations may affect the collaboration relationship between communicating colleagues. They also observed a flexibility of approaches in the communication exchanges as subjects used the voice channel to work towards a viable exchange pattern that would help them solve the problems created by the platform differences. The integrated systems described in the following case studies can also be considered a multiplatform approach to enhance performance and communication patterns when collaborators use heterogeneous computer platforms for their collaboration.

3 Awareware: Audio Windowing Narrowcasting Systems

Vision and audition are the two main human senses for obtaining information about the outside world, and full CSCW applications need both modes (at least!). Visual windowing systems allow multiple and multiwindow applications to share display resources; audio windowing systems, in analogy to graphical windows, can bring order to a cacophony of multiple simultaneous sound sources. Audio windowing can be thought of as a frontend, or articulated user interface, to a system with a spatial sound backend [19, 20, 18, 16]. Using our audio windowing system, users will be able to control the spatialized audio (and other realtime media streams) of inevitable multiparty chatspaces, using the cocktail party effect [5] as well as narrowcasting to make useful sense of the cacophony, as imagined by Figure 4.

Audio windowing narrowcasting commands control superposition of soundscapes. Using the awareness parlance of Benford et al., an aura delimited by a graphical window is like a room, sink attributes affect focus, and source attributes affect nimbus. On a logical level, sound sources and sinks are resources assigned to users. Shared virtual environments like chatspaces require generalized control of user-dependent media streams.

We present two case studies of "awareware," describing audio windowing interfaces for workstations and mobile devices [25], both supporting multiple spaces [22] to enhance narrowcasting conferencing capabilities in CVEs: a workstation WIMP style (**w**indows/**i**con/**m**enu/**p**ointer) GUI (**g**raphical **u**ser **i**nterface), and a MIDlet (**m**obile **i**nformation **d**evices applet) for 2.5 and 3rd-generation mobile phones. The workstation- and mobile-based interfaces encourage use of multiple spaces, leveraging multipresence-enabled conferencing features. Dynamic deletion and creation of avatars controlled by clipboard operations enable teleporting (`cut`&`paste`) and cloning (`copy`&`paste`) avatars in distributed applications and heterogeneous sessions for both interfaces [26], conveying narrowcasting attributes across multiple spaces.

3.1 "Multiplicity": Java3D Workstation-Platformed Multiperspective Interface

The workstation-based audio windowing narrowcasting system, named "Multiplicity" (with a nod to the eponymous movie [46]), developed with JSE and Java3D, runs on Mac OS X, Microsoft Windows, and Sun Solaris. An arbitrary number of avatars can be instantiated and associated with users at runtime. Attributes of narrowcasting functions extend the figurative avatars to denote the invoked filters. Multiplicity can display multiple perspectives from various standpoints, including exocentrically from various strategically placed cameras and egocentrically (both endocentric & tethered) with respect to a selected avatar, in hybrid visual configurations or stereographically.

3.1.1 Multiple Spaces

Multiple spaces are supported in Multiplicity via launching multiple instances of the application. An arbitrary number of applications can be run, corresponding, for example, to domestic, academic, professional, and musical spaces. Upon launching an instance of the application, a user provides a space name, upon which are based all the networked attributes of that space, including channel IDs, position parameters, and narrowcasting attributes. This approach allow users to seamlessly inhibit an arbitrary number of virtual spaces.

3.1.2 Visual Representation of Narrowcasting Operations

A human user can be represented in virtual space by one or more avatars. A figurative avatar in virtual space is naturally humanoid, including especially a head, since it not only embodies a center of consciousness, but also important communication organs: ears, mouth, and eyes. Exclude and include source and sink operations can be visually represented by attributes which can distinguish between operations reflexive, invoked by user associated with a respective avatar, and transitive, invoked by another user in the shared environment.

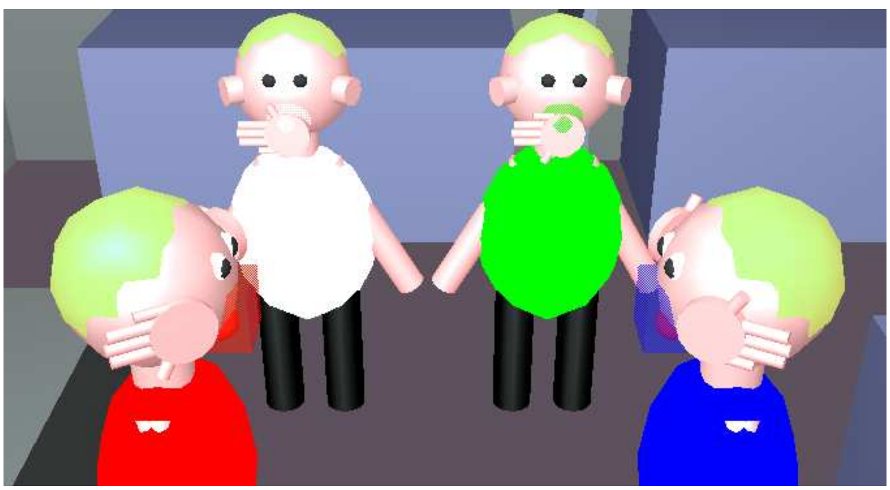

Fig. 6 Exclude narrowcasting operations in Multiplicity: In this example, avatars in the center rear are `mute`d, by `self` (thumb up) and other (thumb down), respectively, while left- and right-most (in the front) avatars are `deafen`ed, by `self` (thumbs down) and other (thumbs up), respectively.

Figurative representations of narrowcasting operations suggest sender- and receiver-side filtering. For exclude operations, virtual hands cover avatars' ears and mouths, with orientation suggesting the nature of the blocking. Exclude audio operations `mute` and `deafen` are shown in Figure 6. A source representing an avatar denotes

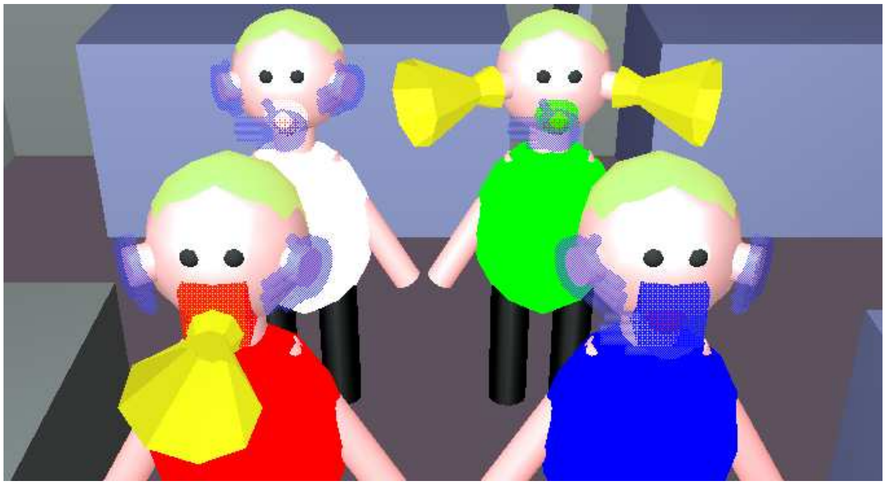

Fig. 7 Include narrowcasting operations in Multiplicity: The avatar in the front left is `select`ed, so its complement (comprising all the other avatars) is `mute`d (denoted by the translucent hands before the mouths), and the avatar in the back right is `attend`ed, so its complement is `deafen`ed (denoted by translucent hands before the mouth).

`mute`dness with an virtual hand clapped over its mouth, oriented differently (thumb up or down) depending on whether the source was `mute`d by its owner (or one of its owners) or another, unassociated user. Hands clapped over the ears are also oriented differently depending on the agent of deafness, thumbs down in the case of reflexive invocation by a user desiring quiet, and thumbs up to denote other-imposed deafness, invoked by another desiring secrecy.

For include operations, `select` and `attend` attributes are denoted by characteristic features, as shown in Figure 7. A megaphone appears in front of `select`ed avatars' mouths, and ear trumpets straddle `attend`ed avatars' ears. If any avatar has been `select`ed, non-`select`ed avatars of the same `self`/non-`self` class are implicitly `mute`d. In the same manner, if any `attend`ed avatars exist in a given space, non-`attend`ed avatars are implicitly `deafen`ed. Translucent hands represent these effects, implicit `mute` represented by a translucent hand clapped over the mouth, and implicit `deafen` represented by translucent hands clapped over the ears. Such narrowcasting attributes are conveyed by avatars even as they move or replicate (via the clipboard) to other spaces, as illustrated by Figure 8.

3.1.3 Visual Representation of Autofocus Operations

In the absence of an autofocus function, a multipresent user associated with multiple sinks might hear each source from several locations, each a manifestation of its respective displacement from each sink. An autofocus function discovers a unique, most sensitive sink for each source by compiling narrowcasting attributes from each

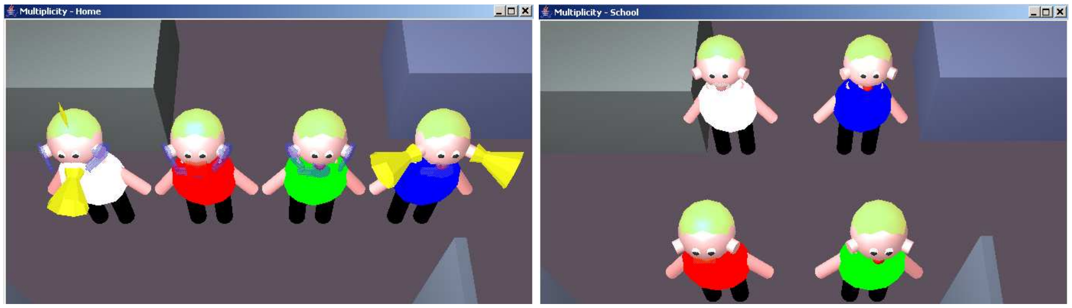

(a) Before clipboard operations (both instances of Multiplicity run in the same login on a computer share the session clipboard): Avatar #0 (left) is `select`ed and #3 (right) is `attend`ed in the Home space, so their complement are implicitly `mut`ed and `deafen`ed, respectively. (No narrowcasting attributes are yet applied to avatars in the School space.)

(b) After teleporting (`cut/paste`) operations: Avatars #0 and #3 have been `cut` from the Home space (along with their narrowcasting attributes, including `self`ness) and `past`ed into the School space, as avatars #4 (midground left) and #5 (midground right). Newly `past`ed avatar #4 is `select`ed and #5 is `attend`ed, so their complements are again implicitly `mut`ed and `deafen`ed, respectively.

(c) After cloning (`copy/paste`) operations: Avatars #0 and #3 have been copied in the Home space (along with their narrowcasting attributes) and `past`ed into the School space. Newly `past`ed avatar #4 is `select`ed and #5 is `attend`ed, so their complements are once again implicitly `mut`ed and `deafen`ed.

Fig. 8 Teleporting and cloning across Home and School instances of Multiplicity

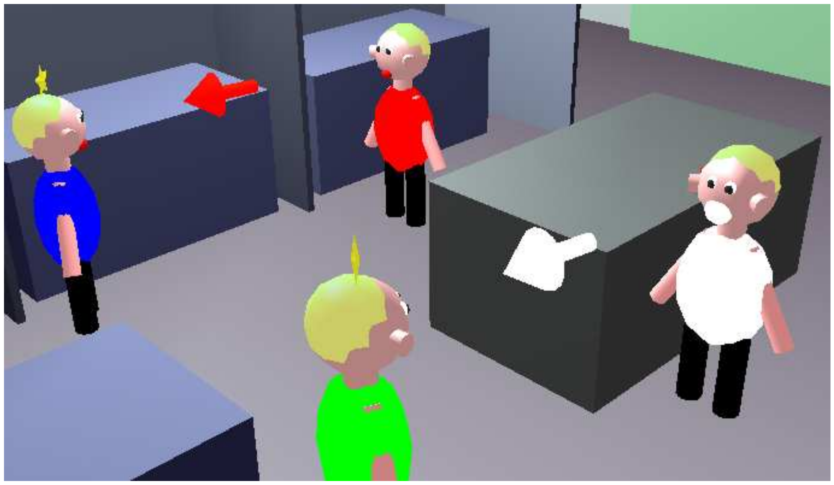

Fig. 9 Autofocus visualization in Multiplicity: Anycasting source → sink vectors are visualized by arrows, flying from each source to its respective "best sink." (The user's sinks, designated as `self`, have stars rotating above their heads.)

source to each sink, and then choosing the respectively closest active ones. (Sink sensitivities depend upon distances from the sources, which are assumed to be omnidirectional.) The position of each source's best sink is denoted by flying animated arrows, as shown in Figure 9.

3.2 "ι·Con": (iαppli DoJa) Mobile Device-Platformed Dynamic Map

A mobile-based audio windowing system, named "ι·Con," developed with Java ME (Micro Edition) and DoJa (DoCoMo Java), runs on (NTT DoCoMo) iαppli mobile phones. Featuring selectable icons with one rotational and two translational degrees of freedom, the "2.5D" dynamic map interface is used to control position, sensitivity, and audibility of avatars in a groupware session. Its isosceles triangle icons are representations of symbolic heads in an orthographic projection, including narrowcasting attributes. The interface also has musical and vibrational cues, to signal mode changes and successful transmission/reception (which feedback is important in wireless communication, as it is much less deterministic than wireline systems).

3.2.1 Multiple Spaces

Multiple spaces are supported directly within the mobile-based interface (there being no symmetric multitasking on the mobile phone operating system), integrated with other applications through a servent (**serv**er/cli**ent** hybrid) HTTP↔TCP/IP gateway. Area-division multiplexing of the graphical display is used for the mobile interface to display multiple spaces. In consideration of the small screen display of mobile devices, the mobile interface currently supports only two virtual rooms ("Home" and "School"), but a general multiwindowing system would allow an arbitrary number of spaces.

3.2.2 Visual Representation of Narrowcasting Operations

Symbolic representations of narrowcasting operations were developed for mobile interface by flattening flattening figurative 3D avatars to 2.5D icons, as seen in Figure 10. In the *ι*·Con application, narrowcasting attributes' graphical displays are triply encoded— by position (before the "mouth" for `mute` and `select`, straddling the "ears" for `deafen` and `attend`), symbol ('+' for include & '–' for exclude), and color (green for assert & red, yellow, and orange for inhibit— by self, other, and implicitly, respectively).

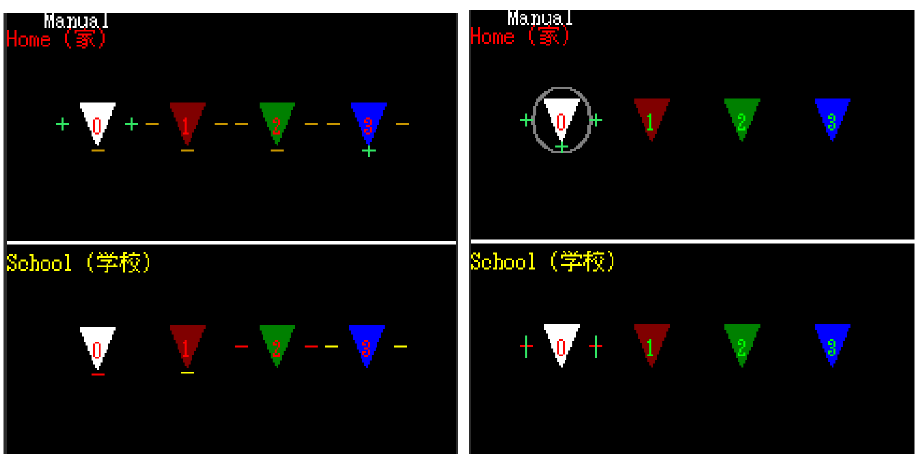

(a) In the Home space, avatar #0 is `attend`ed, so its complement (comprising all the other avatars) is `deafen`ed; and avatar #3 is `select`ed, so its complement is implicitly `mute`d. In the School space, avatars #0 & #2 are respectively `mute`d and `deafen`ed by `self` while avatars #1 & #3 are respectively `mute`d and `deafen`ed by others.

(b) In the Home space, #0 is simultaneously `attend`ed and `select`ed and also selected for rotation (as indicated by its "halo"), while in the School space, #0 is simultaneously `attend`ed and `deafen`ed.

Fig. 10 Narrowcasting attributes on mobile graphical display

3.2.3 Visual Representation of Autofocus Operations

The autofocus function described earlier (in § 1.5 and Figure 9) is also applied to the ι·Con interface, which discovers the best sink for each source (the one for which the source is loudest). A disk is drawn above each source, colored the same as the respective best sink. In the simple case, when narrowcasting attributes are not applied, the ι·Con interface discovers the best sink for each source considering only distance, as shown in Figure 11.

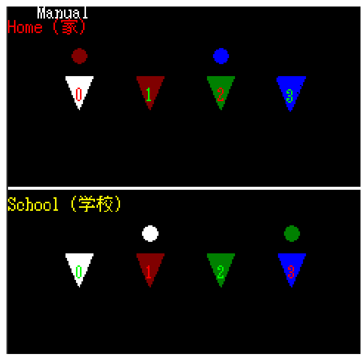

Fig. 11 Autofocus visualization in ι·Con: Home's #1 & #3 and School's #0 & #2 are `self`-designated icons, and therefore candidates for "best sink." In the Home space, #0's & #2's best sinks are #1 & #3 respectively— a circle the same color as #1 is drawn above source #0 while a circle colored the same as #3 is drawn above source #2. In the School space, #1's & #3's best sinks are #0 & #2— a circle the same color as #0 is drawn above source #1, while a circle colored the same as #2 is drawn above source #3.

When narrowcasting attributes are applied to icons in this mobile interface, the best sink for each source depends on the those attributes as well as designation (`self` or non-`self`) of the icons. For example, when a user `selects` a `self`-designated avatar, other `self`-designated avatars are implicitly `mute`d. In the same manner, when a user `selects` a non-`self`-designated avatar, other non-`self`-designated avatars are implicitly `mute`d. Depending on the situation, users can change the determination of the best sink for each source using narrowcasting attributes, as shown in Figure 12(a).

4 Narrowcasting Attributes for Presence Awareness

The configurations described by Benford et al. are re-presented in an original taxonomy [25] shown in Table 2. There are many ways of mapping those situations into arrangements supported by the narrowcasting idioms described in this chapter. For instance, direct analogies between nimbus and source "visibility" (audibility, etc.) and between focus and sink attention allows the equivalences illustrated by Ta-

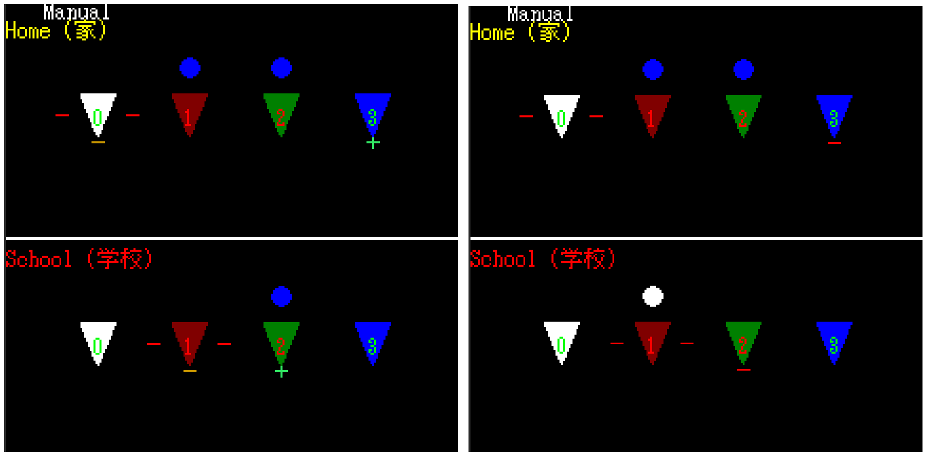

(a) In the Home space, `self`-designated avatars are `deafen`ed (#0) and `select`ed (#3), causing #0 to be implicitly `mute`d. In the School space, non-`self`-designated avatars are `deafen`ed (#1) and `select`ed (#2), causing #1 to be implicity muted.

(b) In the Home space, `self`-designated avatars are `deafen`ed (#0) and `mute`d (#3). In the School space, non-`self`-designated avatars are `deafen`ed (#1) and `mute`d (#2).

Fig. 12 Narrowcating attributes are applied to `self` and non-`self` avatars and autofocus behavior displayed in different situations. Home #0 & #3 and School #0 & #3 are `self`-designated avatars.

ble 3, as `include` narrowcasting operations (`attend`/`select`), and Table 4, as `exclude` narrowcasting operations (`deafen`/`mute`). A subject focuses attention on an object as a sink focuses on a source. An object attracts attention from a subject as the autofocus function causes an "anycasting" source to discover a best sink. For example, if a source is `mute`d, either by its owner or the other participant, its nimbus excludes the other avatar. For instance, one might hold their hand over their mouthpiece (microphone) or push a "Hold" button (like that in audio chatspaces) of a phone handset to block the transmission, or use a "sneeze" button to freeze a video stream.

As a narrowcasting interface is designed for more than two participants, there are differences between, for instance, `select`ing a source and `mut`ing its compliment. Tables 3 & 4 show a coarse projection of a much more complicated space. Neither do Tables 2–4 distinguish between narrowcasting attributes invoked by oneself vs. by another. That is, capability can be determined by combined narrowcasting attributes, independent of agent, but in context such distinctions are very important. A can't hear B if B is `mute`d by A or by B himself, but there is a big difference socially, especially in presence of a third actor C, who could hear B `mute`d transitively by A or others (besides C) but not reflexively (by B). If all the attributes invoked by

Table 2 Modes of mutual awareness (Benford et al.). Circles depict the nimbus projected by an object, and arrows depict the direction of the subject's focus. Because of symmetry, the relation is basically reducible to a triangular matrix, with analogous transposition reflected across the main diagonal.

		B focused on A		B not focused on A	
		A ∈ B's nimbus	A ∉ B's nimbus	A ∈ B's nimbus	A ∉ B's nimbus
A focused on B	B ∈ A's nimbus	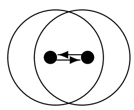 10. fully reciprocal mutual awareness	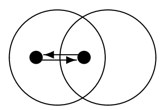 9. withdrawal	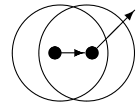 7. monitoring	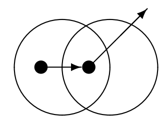 5. ignoring
	B ∉ A's nimbus	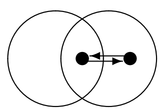 (9.)	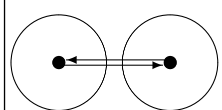 8. mutual minimal awareness	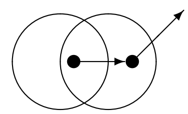 6. eavesdropping	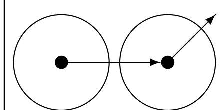 4. minimal asymmetrical awareness
A not focused on B	B ∈ A's nimbus	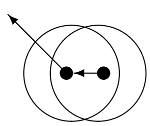 (7.)	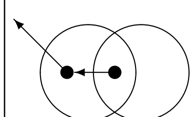 (6.)	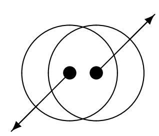 3. mutual overhearing	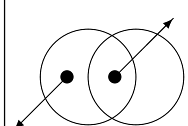 2. overhearing / distraction
	B ∉ A's nimbus	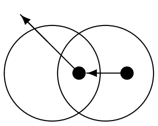 (5.)	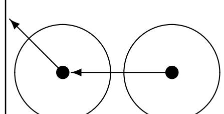 (4.)	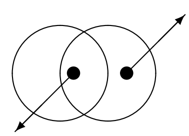 (2.)	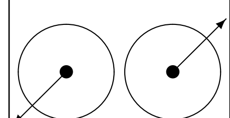 1. no mutual awareness

Table 3 `Include` narrowcasting modes for mutual awareness (`attend` and `select`).

		B `attend`ed		B not `attend`ed	
		B `select`ed by A	B not `select`ed by A	B `select`ed by A	B not `select`ed by A
A `attend`ed	A `select`ed by B	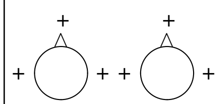 10. fully reciprocal mutual awareness	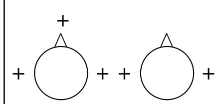 9. asymmetric reciprocal mutual awareness	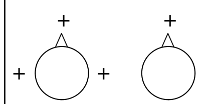 7. asymmetric reciprocal mutual awareness	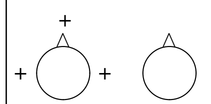 5. asymmetric awareness
	A not `select`ed by B	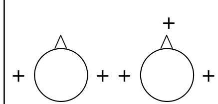 (9.)	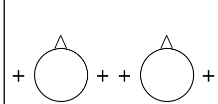 8. mutual overhearing	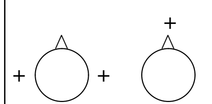 6. lurking, asymmetric awareness	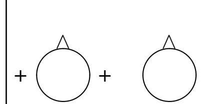 4. minimal asymmetric awareness
A not `attend`ed	A `select`ed by B	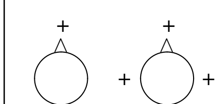 (7.)	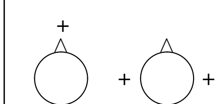 (6.)	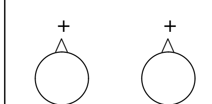 3. symmetric awareness	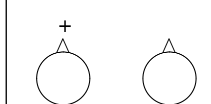 2. minimal asymmetric awareness
	A not `select`ed by B	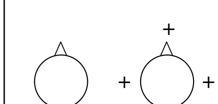 (5.)	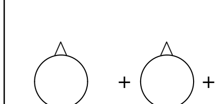 (4.)	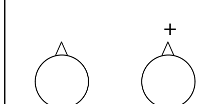 (2.)	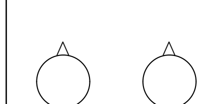 1. minimal mutual awareness

Table 4 `Exclude` narrowcasting modes for mutual awareness (`deafen` and `mute`).

		B not `deafen`ed		B `deafen`ed	
		B not `mute`d	B `mute`d	B not `mute`d	B `mute`d
A not `deafen`ed	A not `mute`d	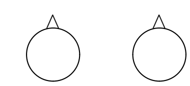 10. minimal mutual awareness	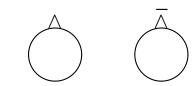 9. censor, withdrawal	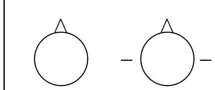 7. monitoring	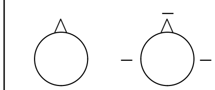 5. isolate
	A `mute`d	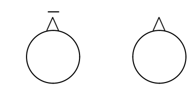 (9.)	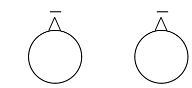 3. null symmetric awareness	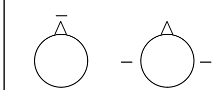 6. eavesdropping	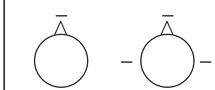 4. null asymmetric awareness
A `deafen`ed	A not `mute`d	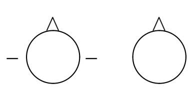 (7.)	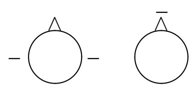 (6.)	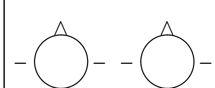 8. null symmetric awareness	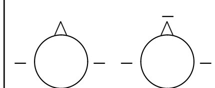 2. null asymmetric awareness
	A `mute`d	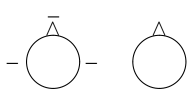 (5.)	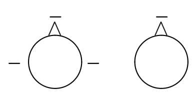 (4.)	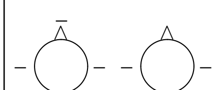 (2.)	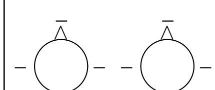 1. no mutual awareness

Table 5 Corresponding narrowcasting modes for mutual awareness (`attend` and `mute`). The +s at the ears, straddling the iconic heads, denote explicitly enabled sinks, and −s before the mouths denote disabled sources.

		B `attended`		B not `attended`	
		B not `muted`	B `muted`	B not `muted`	B `muted`
A `attended`	A not `muted`	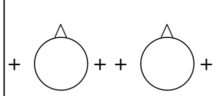 10. fully reciprocal mutual awareness	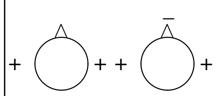 9. withdrawal	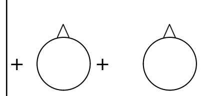 7. monitoring	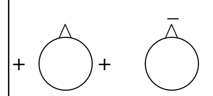 5. ignoring
	A `muted`	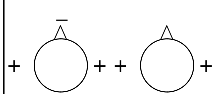 (9.)	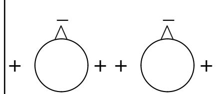 8. mutual minimal awareness	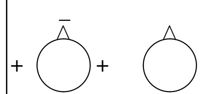 6. eavesdropping	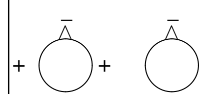 4. minimal asymmetrical awareness
A not `attended`	A not `muted`	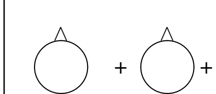 (7.)	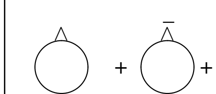 (6.)	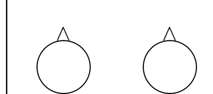 3. mutual overhearing	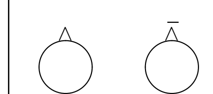 2. overhearing / distraction
	A `muted`	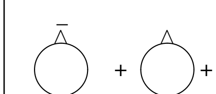 (5.)	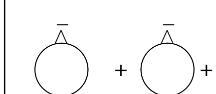 (4.)	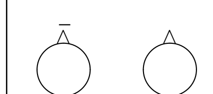 (2.)	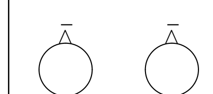 1. no mutual awareness

Table 6 Corresponding narrowcasting modes for mutual awareness (`deafen` and `select`). The −s at the ears, straddling the iconic heads, denote disabled sinks, and +s before the mouths denote enabled sources.

		B not `deafen`ed		B `deafen`ed	
		B `select`ed by A	B not `select`ed by A	B `select`ed by A	B not `select`ed by A
A not `deafen`ed	A `select`ed by B	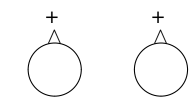 10. fully reciprocal mutual awareness	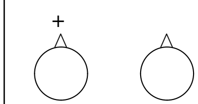 (7.)	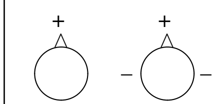 (9.)	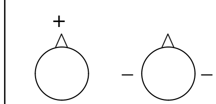 (5.)
	A not `select`ed by B	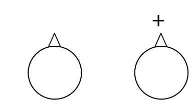 7. monitoring	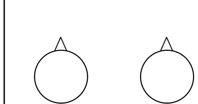 3. mutual overhearing	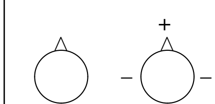 6. eavesdropping	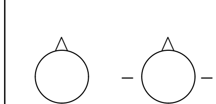 (2.)
A `deafen`ed	A `select`ed by B	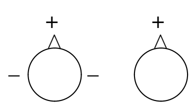 9. withdrawal	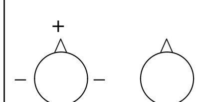 (6.)	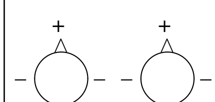 8. mutual minimal awareness	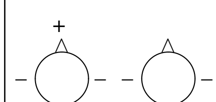 (4.)
	A not `select`ed by B	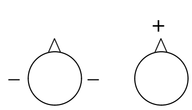 5. ignoring	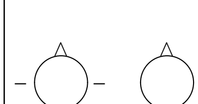 2. overhearing/distraction	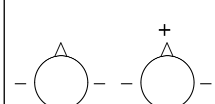 4. minimal asymmetrical awareness	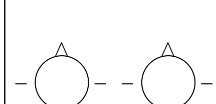 1. no mutual awareness

one's self vs. by another only for Table 4 are considered, there are many possible situations, some of them symmetric. For example, if A `mute`s himself, nobody can hear A (except A), so it doesn't matter if A is also `mute`d by B.

5 Future Research

5.1 Role-Based Issues

Roles are a powerful concept for facilitating distributed systems management and enforcing access control [49, 42, 36, 54]. The basic idea of role-based collaboration is that a collaborative system that can designate explicity what objects users can access with which specific rights, and can also designate which users they can manage or communicate with, they can then accomplish their jobs meaningfully and efficiently.

Should a student be allowed `deafen` a teacher, or a teenager be allowed to `mute` a parent? The models described in this chapter ignore such higher-order considerations, like visibility of applied attributes. The interfaces described are transparent: any attributes invoked by any participant are revealed to the other users in the session. Such "perfect information" (from game study, in which all actors have access to all information) begs the question: If A `mute`s B, should B always be aware of it? Depending on the conditions, such transparency could be appropriate or not. A parent might insist upon the ability to override a teenager's 'ignore' command: "How dare you mute me‽" Such role-based issues are subtle and sociological, and are the subject of ongoing consideration.

5.2 Next-Generation Mobile Phones

4G mobile services [37] will include network [34] technology integration, SDR (**s**oftware-**d**efined **r**adio), and advanced multimedia [53] mobile communications (IPv6, high-resolution video transmission digital broadcasting, security, etc.) including 3D VR-style interfaces. The catchphrase for 4G is ABC: "**a**lways **b**est **c**onnected," suggesting possibility of persistent sessions, as imagined by Figure 13.

5.3 Convergence

Besides wireline-connected workstation-based interfaces, narrowcasting might find an even more fertile platform in mobile devices [33]. The '4-play' convergence of telephony, television/video, internet, and wireless is driving a proliferation of new

"I have to hang up now. You just walked through the door."

Fig. 13 "**A**lways **b**est **c**onnected" yields practically persistent sessions.

devices and services. Mobile terminals, almost as intimate as clothing, are a kind of wearable computer, and a diversity of ever-next-generation functionalities and form factors for smartphones is emerging, including mobile stereotelephony, inspired by cyberspatial audio [17] and augmented audio models. Meanwhile, location-based services— along with seamless handoff, FMC (fixed-mobile convergence), and heterogeneous roaming via MIMO (multiple input/multiple output) smart antennas leading to software-defined radio (SDR) and cognitive radio— leverage geolocation and portable GPS/GIS.

5.4 "Polite Calling" for Social Gracefullness

Advanced sensing— including optical systems, position trackers, and motion sensors in mobile phones— encourages ubicomp (ubiquitous computing) and ambient intelligence, including an indirect awareness of user status and availability, "presence," which, along with explicit status settings by a user, enable "polite calling" that is respectful of the accessibility of a callee, including distractedness or preoccupation, sleep, social context, etc. Agents delegated on behalf of a caller and callee could negotiate an appropriate interruption, based upon caller insistence and callee receptiveness, including dropping down to voice- or video-mail, or ringing through a "don't disturb" in the case of emergencies.

Such filters will also increase the number of welcome calls, since callers will no longer have to self-censor, secure in the knowledge that a callee is appropriately shielded. (For simple example, one now hesitates to call even intimate friends and relations in the middle of the night for fear of waking them up, but if it was assumed that everyone had a gateway for active call screening that knew about their sleeping, then one could call spontaneously, hoping that the callee was receptive, but confident that if they weren't that "the machine" would intercept and not bother the sleeping callee for non-urgent matters.)

Like some proposed measures to deal with telemarketers' SPIT (spam over IP telephony) and SPIM (spam over instant messaging), perhaps some negotiations would force the caller to electronically post a financial surety or bond of indemnification into escrow, which the callee could keep if the call is deemed a nuisance or not sufficiently important.

The distinction will blur between "calling" someone, to establish a circuit, and "calling to" someone, to get their attention. Articulated models of privacy like narrowcasting will allow users to distribute their attention, availability, and virtual presence. Multipresence and persistent channels, encouraged by ABC (always best connected) networks, will extend the way people communicate.

6 Conclusion

The basic goal of the research described in this chapter is to develop idioms for selective attention, privacy, and presence: narrowcasting for groupware applications, whether the interface is via workstation or a nomadic device like a mobile phone. We described deployment a multiplatform implementation of multipresence-enabled narrowcasting functions, including autofocus determination for both workstations and mobile devices. The workstation application features a multiperspective interface, including logical separation of eyes and ears (virtual camera and stereo microphones), exploiting the "phantom source" feature we developed. The mobile interface features equivalent exocentric narrowcasting commands, displayed and controlled in a manner appropriate for the unique form factor of the contemporary mobile phone. The platform-agnostic deployment of the audio narrowcasting idioms— including `deafen`, `mute`, `select`, and `attend`— encourages modernization of office- and mobile-based conferencing, leveraging session integration across coextensive spaces and anticipating multipresence enabled by higher bandwidth and more durable or even persistent mobile connectivity. One will have presence in many different places as well as ability to shift attention back and forth. For instance, one's family members, schoolmates, friends, etc. will have virtual copresence and one can virtually go back and forth among different spaces.

Normally, what one sees is tightly aligned with what one hears, since the eyes and ears are "concentric," locked together as they are in one's head, but users can fork

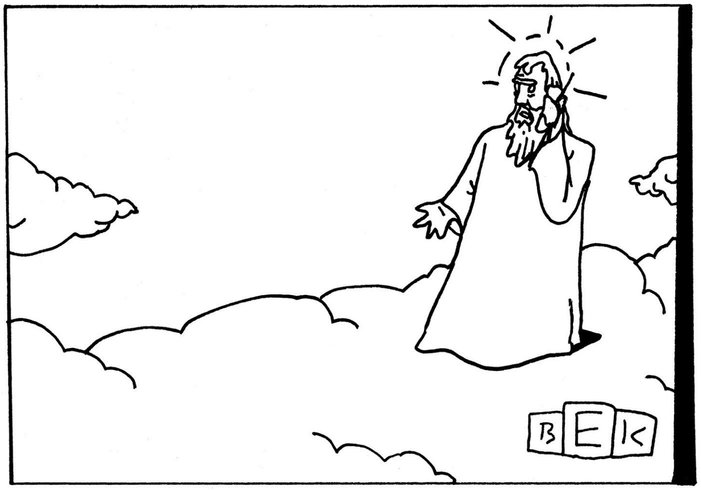

"Sorry, I can't—I have to be everywhere."

Fig. 14 Divine ubiquity. (© The New Yorker Collection 2003 Bruce Eric Kaplan from `cartoonbank.com`. All rights reserved.)

themselves through designation of multiple avatars, compositing phantom sources via the superposition of multiple sinks' soundscapes. For instance, one might "fork presence" in virtual rooms corresponding to home (chatspace), school (teleconference), and music (virtual concert). Activity or information in a space might cause the user to focus on that particular soundscape, using narrowcasting functions [43]. As suggested by Figure 14, being anywhere is better than being everywhere, since it is selective; multipresence is distilled ubiquity, narrowcasting-enabled audition (for sinks) or address (for sources) of multiple objects of regard. This research can be considered an extension of presence technology [56], and anticipates deployment of such narrowcasting protocols into session protocols like SIP/SIMPLE [35] [11] [3, 2, 4] or the internet infrastructure (routers, etc.) itself.

References

1. Ackerman, M.S.: Privacy in pervasive environments: next generation labeling protocols. **P**ersonal **U**biquitous **C**omput. **8**(6), 430–439 (2004). DOI http://dx.doi.org/10.1007/s00779-004-0305-8
2. Alam, M.S., Cohen, M., Ahmed, A.: Articulated Narrowcasting for Privacy and Awareness in Multimedia Conferencing Systems and Design for Implementation Within a SIP Framework. JVRB: **J**. of **V**irtual **R**eality and **B**roadcasting **4**(9) (2007)
3. Alam, M.S., Cohen, M., Ahmed, A.: Narrowcasting— Controlling Media Policy in SIP Multimedia Conferencing. In: Proc. IEEE CCNC: 4th **C**onsumer **C**ommunications and **N**etworking **C**onf. Las Vegas (2007)

4. Alam, S., Cohen, M., Ashir, A., Villegas, J.: Narrowcasting in SIP: Articulated Privacy Control. In: A. Syed, M. Ilyas (eds.) SIP Handbook: Services, Technologies, and Security of Session Initiation Protocol, chap. 14, pp. 323–345. Taylor and Francis (2008)
5. Aoki, P.M., Romaine, M., Szymanski, M.H., Thornton, J.D., Wilson, D., Woodruff, A.: The Mad Hatter's Cocktail Party: a social mobile audio space supporting multiple simultaneous conversations. In: Proc. CHI: ACM Conf. on **C**omputer-**H**uman **I**nteraction, pp. 425–432. Ft. Lauderdale (2003)
6. Benford, S., Bowers, J., Fahlén, L.E., Greenhalgh, C., Snowdon, D.: User embodiment in collaborative virtual environments. In: Proc. CHI: SIGCHI Conf. on Human factors in computing systems, pp. 242–249. ACM Press/Addison-Wesley Publishing Co., Denver (1995). ISBN 0-201-84705-1
7. Benford, S., Bowers, J., Fahlén, L.E., Mariani, J.A., Rodden, T.: Supporting Cooperative Work in Virtual Environments. Comput. J. **37**(8), 653–668 (1994)
8. Benford, S., Greenhalgh, C., Reynard, G., Brown, C., Koleva, B.: Understanding and constructing shared spaces with mixed-reality boundaries. ACM Trans. Comput.-Hum. Interact. **5**(3), 185–223 (1998). DOI http://doi.acm.org/10.1145/292834.292836
9. Benford, S., Greenhalgh, C., Rodden, T., Pycock, J.: Collaborative virtual environments. Communications of the ACM **44**(7), 79–85 (2001). DOI http://doi.acm.org/10.1145/379300.379322
10. Blauert, J.: Spatial Hearing: The Psychophysics of Human Sound Localization, revised edn. MIT Press (1997). ISBN 0-262-02413-6
11. Boyer, D., Ginsberg, A., Goud, V., Handel, M.: Presence awareness for future telecommuniction systems. In: A. Pakštas, R. Komiya (eds.) Virtual Reality Technologies for Future Telecommunications Systems, chap. 4, pp. 31–44. Wiley, West Sussex, England (2002). ISBN 0-470-84886-3
12. Christein, H., Schulthess, P.: A General Purpose Model for Presence Awareness. In: Proc. Fourth Int. Conf. on Distributed Communities on the Web. Sydney (2002)
13. Cockburn, A., Weir, P.: An investigation of groupware support for collaborative awareness through distortion-oriented views. IJHCI: **I**nt. **J**. of **H**uman-**C**omputer **I**nteraction **11**(3), 231–255 (1999)
14. Cohen, M.: Quantity of presence: Beyond person, number, and pronouns. In: T.L. Kunii, A. Luciani (eds.) Cyberworlds, chap. 19, pp. 289–308. Springer-Verlag, Tokyo (1998). ISBN 4-431-70207-5; `www.u-aizu.ac.jp/~mcohen/welcome/publications/QuantityOfPresence.pdf`
15. Cohen, M.: Exclude and include for audio sources and sinks: Analogs of mute & solo are deafen & attend. Presence: Teleoperators and Virtual Environments **9**(1), 84–96 (2000). ISSN 1054-7460; `www.u-aizu.ac.jp/~mcohen/welcome/publications/ie1.pdf`
16. Cohen, M.: Articulated modeling of distributed privacy: Transitive closure of composition of narrowcasting and multipresence. In: Proc. CIT: Sixth Int. Conf. on **C**omputer and **I**nformation **T**echnology. Seoul (2006)
17. Cohen, M., Herder, J., Martens, W.L.: Cyberspatial Audio Technology. **J**. **A**cous. **S**oc. **J**ap. (English) **20**(6), 389–395 (1999). ISSN 0388-2861; `www.u-aizu.ac.jp/~mcohen/welcome/publications/JASJ-reviewE.pdf`
18. Cohen, M., Koizumi, N.: Virtual gain for audio windows. Presence: Teleoperators and Virtual Environments **7**(1), 53–66 (1998). ISSN 1054-7460
19. Cohen, M., Ludwig, L.F.: Multidimensional audio window management. IJMMS: J. of Person-Computer Interaction **34**(3), 319–336 (1991). Special Issue on Computer Supported Cooperative Work and Groupware. ISSN 0020-7373
20. Cohen, M., Ludwig, L.F.: Multidimensional audio window management. In: S. Greenberg (ed.) Computer Supported Cooperative Work and Groupware, chap. 10, pp. 193–210. Academic Press, London (1991). ISBN 0-12-299220-2
21. DiPaola, S., Collins, D.: A 3D Virtual Environment for Social Telepresence. In: Proc. Western Computer Graphics Symp. Vernon, BC; Canada (2002)

22. Dix, A., Friday, A., Koleva, B., Rodden, T., Muller, H., Randell, C., Steed, A.: Managing multiple spaces. In: P. Turner, E. Davenport (eds.) Spaces, Spatiality and Technology, *The Computer Supported Cooperative Work Series*, vol. 5. Springer (2005). ISBN 1-4020-3272-2
23. Donath, J.S., Viégas, F.B.: Chat circles series: Explorations in designing abstract graphical communication interfaces. In: Proc. DIS, Conf. on Designing Interactive Systems. London, England (2002)
24. Dourish, P., Bly, S.: Portholes: Supporting awareness in a distributed work group. In: Proc. CHI'92: Human Factors in Computing Systems, pp. 541–547. Addison-Wesley (1992)
25. Fernando, O.N.N., Adachi, K., Duminduwardena, U., Kawaguchi, M., Cohen, M.: Audio Narrowcasting and Privacy for Multipresent Avatars on Workstations and Mobile Phones. IEICE Trans. on Information and Systems **E89-D**(1), 73–87 (2006). `ietisy.oxfordjournals.org/cgi/content/refs/E89-D/1/73`, ISSN 0916-8532
26. Fernando, O.N.N., Saito, G., Duminduwardena, U., Tanno, Y., Cohen, M.: Cloning and teleporting avatars across workstations and mobile devices in collaborative virtual environments: Clipboard operations for virtual reality. In: Proc. ICIA'05: **I**nt. **C**onf. on **I**nformation and **A**utomation. Colombo, Sri Lanka (2005). `www.ent.mrt.ac.lk/iml/ICIA2005/Papers/JP001CRC.pdf`
27. Gilkey, R.H., Anderson, T.R. (eds.): Binaural and Spatial Hearing in Real and Virtual Environments. Lawrence Erlbaum Associates, Mahway, NJ (1997). ISBN 0-8058-1654-2
28. Greenhalgh, C., Benford, S.: Massive: A collaborative virtual environment for teleconferencing. ACM Trans. on Computer-Human Interaction **2**(3), 239–261 (1995)
29. Grudin, J.: Computer-supported cooperative work: History and focus. (IEEE) Computer **27**(5), 19–26 (1994). DOI http://doi.ieeecomputersociety.org/10.1109/2.291294
30. Gutwin, C., Greenberg, S.: Design for individuals, design for groups: tradeoffs between power and workspace awareness. In: Proc. CSCW: ACM Conf. on Computer supported cooperative work, pp. 207–216. ACM Press, Seattle (1998). ISBN 1-58113-009-0
31. Gutwin, C., Greenberg, S.: Effects of awareness support on groupware usability. In: Proc. CHI: SIGCHI Conf. on Human Factors in Computing Systems, pp. 511–518. ACM Press/Addison-Wesley Publishing Co., Los Angeles (1998). ISBN 0-201-30987-4
32. Haas, H.: The influence of a single echo on the audibility of speech. J. **A**ud. **E**ng. **S**oc. **20**, 146–159 (1972)
33. Hazas, M., Scott, J., Krumm, J.: Location-aware computing comes of age. IEEE Computer Magazine **37**(2), 95–97 (2004)
34. Huber, J.F.: Mobile next-generation networks. (IEEE) MultiMedia **11**(1), 72–83 (2004). DOI http://doi.ieeecomputersociety.org/10.1109/MMUL.2004.1261110
35. Johnston, A.B.: SIP: Understanding the Session Initiation Protocol. Artech House, London (2004). ISBN 1580531687
36. Kern, A., Kuhlmann, M., Schaad, A., Moffett, J.: Observations on the role life-cycle in the context of enterprise security management. In: Proc. SACMAT'02: 7th ACM Symp. on Access Control Models and Technologies, pp. 43–51. ACM Press, Monterey, CA; USA (2002). ISBN 1-58113-496-7
37. Kobylarz, T.J.: Beyond 3G: Compound wireless services. (IEEE) Computer **37**(9), 23–28 (2004). DOI http://doi.ieeecomputersociety.org/10.1109/MC.2004.120
38. Marmasse, N., Schmandt, C., Spectre, D.: Watchme: Communication and awareness between members of a closely-knit group. In: N. Davies, E.D. Mynatt, I. Siio (eds.) Ubicomp, *Lecture Notes in Computer Science*, vol. 3205, pp. 214–231. Springer (2004)
39. Metaxas, G., Markopoulos, P.: Abstractions of awareness. In: P. Markopoulos, B. de Ruyter, W. Mackay (eds.) Awareness Systems: advances in theory, methodology and design, chap. 6. Springer (2009). Human Computer Interaction
40. Muller, M.J., Wharton, C., McIver Jr., W.J., Laux, L.: Toward an hci research and practice agenda based on human needs and social responsibility. In: Proc. CHI: SIGCHI Conf. on Human Factors in Computing Systems, pp. 155–161. ACM Press, New York (1997). DOI http://doi.acm.org/10.1145/258549.258640
41. Papadopoulos, C.: Improving awareness in mobile cscw. IEEE Trans. on Mobile Computing **5**(10), 1331–1346 (2006)

42. Park, J.S., Sandhu, R., Ahn, G.J.: Role-based access control on the web. ACM Trans. Inf. Syst. Sec. **4**(1), 37–71 (2001). DOI http://doi.acm.org/10.1145/383775.383777
43. Pentland, A.S.: Socially aware computation and communication. (IEEE) Computer **38**(3), 33–40 (2005). DOI http://dx.doi.org/10.1109/MC.2005.104
44. Radenkovic, M., Greenhalgh, C., Benford, S.: Deployment issues for multi-user audio support in CVEs. In: Proc. ACM Symp. on **V**irtual **R**eality **S**oftware and **T**echnology, pp. 179–185. ACM Press, Hong Kong (2002). DOI http://doi.acm.org/10.1145/585740.585770
45. Raento, M., Oulasvirta, A.: Designing for privacy and self-presentation in social awareness. PUC: **P**ersonal and **U**biquitous **C**omputing **12**(7), 527–542 (2008)
46. Ramis, H., Miller, C.: Multiplicity (1996). Columbia Pictures Corp.
47. Rashid, O., Mullins, I., Coulton, P., Edwards, R.: Extending cyberspace: location based games using cellular phones. Comput. Entertain. **4**(1), 4 (2006). DOI http://doi.acm.org/10.1145/1111293.1111302
48. Read, K., Maurer, F.: Developing mobile wireless applications. (IEEE) Internet Computing **07**(1), 81–86 (2003). DOI http://doi.ieeecomputersociety.org/10.1109/MIC.2003.1167345
49. Sandhu, R., Bhamidipati, V., Munawer, Q.: The ARBAC97 model for role-based administration of roles. ACM Trans. Inf. Syst. Secur. **2**(1), 105–135 (1999). DOI http://doi.acm.org/10.1145/300830.300839
50. Schroeder, R.: Beyond presence and copresence: A phenomenological account of experiences in shared virtual environments. In: Proc. Presence. Aalborg, Denmark (2003)
51. Slater, M., Linakis, V., Usoh, M., Kooper, R.: Immersion, Presence, and Performance in Virtual Environments: An Experiment using Tri-Dimensional Chess. In: M. Green (ed.) VRST: ACM **V**irtual **R**eality **S**oftware and **T**echnology, pp. 163–172 (1996). ISBN 0-89791-825-8
52. Slater, M., Steed, A.: Meeting people virtually: Experiments in shared virtual environments. In: R. Schroeder (ed.) The Social Life of Avatars: Presence and Interaction in Shared Virtual Environments, pp. 146–171. Springer, London (2002)
53. Subramanya, S.R., Yi, B.K.: User-controlled, multimedia-enhanced communication using prior knowledge and experience. (IEEE) MultiMedia **12**(2), 90–95 (2005). DOI http://dx.doi.org/10.1109/MMUL.2005.40
54. Tolone, W., Ahn, G.J., Pai, T., Hong, S.P.: Access control in collaborative systems. ACM Comput. Surv. **37**(1), 29–41 (2005). DOI http://doi.acm.org/10.1145/1057977.1057979
55. Tsingos, N.: Perceptual audio rendering of complex virtual environments. In: Proc. of ACM acronymSiggraph 2004, pp. 249–258. San Diego (2004)
56. Vaughan-Nichols, S.J.: Presence technology: More than just instant messaging. (IEEE) Computer **36**(10), 11–13 (2003). ISSN 0018-9162
57. Velez, M., Tremaine, M.M., Sarcevic, A., Dorohonceanu, B., Krebs, A., Marsic, I.: "Who's in charge here?" Communicating Across Unequal Computer Platforms. ACM Trans. Comput.-Hum. Interact. **11**(4), 407–444 (2004). DOI http://doi.acm.org/10.1145/1035575.1035579
58. Viégas, F.B., Donath, J.S.: Chat circles. In: Proc. CHI: ACM Conf. on **C**omputer-**H**uman **I**nteraction, pp. 9–16. Pittsburgh (1999). ISBN 0-201-48559-1; `www.media.mit.edu/~fviegas/chat_circles.pdf`
59. Vyas, D., van de Watering, M.R., Eliëns, A., van der Veer, G.C.: Being social work: Designing for playfully mediated social awareness in work environments. In: A. Venkatesh, T. Gonzalves, A. Monk, K. Buckner (eds.) Home Informatics and Telematics: ICT for the Next Billion. Springer (2007). Vol. 241; IFIP: Int. Federation for Information Processing; ISBN 978-0-387-73696-9
60. Vyas, D., van de Watering, M.R., Eliëns, A., van der Veer, G.C.: Engineering social awareness in work environments. In: Universal Access in Human-Computer Interaction: Proc. 12th Int. Conf. on Human-Computer Interaction, pp. 254–263. Beijing (2007). LNCS 4555; ISBN 978-3-540-73280-8
61. Wallach, H., Newman, E.B., Rosenzweig, M.R.: The precedence effect in sound localization. American J. of Psychology **57**, 315–336 (1949)