A View From the Luminous Room

John Underkoffler

MIT Media Laboratory, Cambridge, USA

Abstract: A new body of research, concerning the divergence of information systems away from today's ubiquitous notion of 'computer', is emerging. Though the various threads of inquiry diverge significantly in their approach, each generally seeks some intimate integration of computational resources with its user's immediate or extended environment. We introduce here a new project called the *Luminous Room*, which uses dynamically controlled video projection to permit information access and manipulation throughout an architectural space. A description of the working components of the system is followed by an explication of the philosophies that provide both motivation and a conceptual scaffolding for the work. Finally, we offer speculations about the way in which formal architecture and 'environmental' information systems like the *Luminous Room* might co-develop.

Introduction

There is a growing collection of research that eschews the traditional conception of human–computer interaction — this comprising only the ubiquitous monitor, keyboard, and mouse — in favour of transactions carried out on more decidedly human terms. Though diverse in their approaches and implementations, the experiments in this collection are strongly related by their ultimate goals. The most prominent of these is the conviction that the trend of obliging users to adapt or conform to the characteristics and vicissitudes of the digital machine needs to be reversed, placing instead the burden of conformity (to existing human idioms and environments) on the computer. A representative group of such research initiatives was presented in an earlier article [1], which collected them under the aggregate title 'Antisedentary Beigeless Computing'.

The *Luminous Room* project, ongoing at the MIT Media Laboratory, maps a new region of the ABC frontier. This system achieves a particular kind of environmental integration by coupling computer-driven video projection with hard-coded and vision-updated understandings of its surrounding region. Optical and mechanical components give the *Luminous Room* system full projective and visual access to the space in which it is resident; in general the work seeks to render all parts of an architectural expanse information-bearing and the whole an arena for digital interaction.

In the following section we review the existing experimental systems, explored in detail in the prenominate article, that together define the field of Antisedentary Beigeless Computing and that form the background for the work described here. In Section 3 we introduce the *Luminous Room* project itself, discussing the particulars of the system's construction and the extant and nascent applications responsible for the 'services' offered by the system. Then in Section 4 we tender the general philosophy that has given rise to the *Luminous Room* and that suggests its ideal, 'eventual' form. Finally, in Section 5, we analyse what role ABC systems may play in the context of architecture, and how both architecture and ABC systems like the Luminous Room may eventually come to modulate each other.

Context

Several historical experiments, together with a handful of ongoing and recently completed projects (most undertaken in the 1990s), illuminate the boundary of what ABC is and what ABC systems intend. The *Luminous Room* project will be seen at once to share various characteristics with certain of these extant systems and at the same time to define a new region of ABC territory. (Again, a more extensive version of the system descriptions below and the positioning of the Luminous Room system within the ABC taxonomy is available in [1]).

Historical systems

- *ARCMAC.* Nicholas Negroponte's Architecture Machine Group built a dedicated information environment called the 'Media

Personal Technologies (1997) 1:49-59

Room'. This room was outfitted with all manner of input – output equipment, from touchscreens and videodisk players to eye-trackers, voice recognition hardware, and Polhemus sensors, and served as a testbed platform for several experimental systems [2]. Two of the most relevant experiments are these:

- *Put-That-There* used only natural utterances and direct pointing, these smoothly integrated and interdependently analysed, to allow a user to create, edit, and manipulate geometric shapes and symbolic icons in a cartographic context on a large projection wall [3].
- *World of Windows* presented a user with a wall full of small 'windows' displaying text, slide-shows, static images, and dynamic 'movies', the whole navigated through gaze and visual attention [4].
- *VIDEOPLACE*. The several-decades-long continuum of Myron Krueger's work, beginning with *METAPLAY* and *PSYCHIC SPACE* and continuing through *VIDEOPLACE* and its contemporaries *VIDEODESK* and *VIDEOTOUCH*, is an important precedent for environmentally-dispersed, direct-manipulation interfaces. These systems attribute meaning to coloured silhouette images of each user displayed via real-scale video projection, providing an interaction mechanism based solely on the positions, configurations, and motions of these human 'shadows' [5].

Present-day systems

- *Digital Desk*. This is Pierre Wellner's hybrid real – virtual system in which a normal desk surface is augmented overhead with a video projector and a video camera. Paper documents on the desk's surface are intermingled with digital documents, and the system provides mechanisms for moving seamlessly between the two forms [6].
- *KARMA*. Steven Feiner and his colleagues use a 'head-up' display and position-tracking hardware to overlay dynamically-evolving graphical information onto portions of a laboratory environment. In one application, an automatic illustration design and layout system generates spatially-aligned adjustment and maintenance instructions for a laser printer, these graphical instructions appearing literally atop the printer [7].
- *Ubiquitous Computing*. Mark Weiser, pursuing the notion that computational resources should be literally ubiquitous throughout a workspace 'like physical paper', has implemented representative display-and-input devices at three scales, each generally comprising an LCD display, an electronic stylus, and an IR wireless link. Fluid intercommunication among these distributed components allows (for example) a communal sketching system [8].
- *metaDESK*. This project (and its predecessor the *Bricks* system) is Hiroshi Ishii's migration of familiar symbolic tools back from a virtual pixels-only representation toward physical instantiation in the real world. The *metaDESK* employs a large bottom-projected workspace work and reconceives familiar GUI elements (windows, icons, menus) as TUI (Tangible User Interface) elements, permitting for example the placement of small physical models of recognisable buildings to control the scrolling, scaling, and rotation of a map [9].
- *The Spatially Aware Palmtop*. George Fitzmaurice demonstrates a system that is one plausible model for information organisation in a milieu of small, plentiful, and portable display-and-input devices. Here a palmtop unit, outfitted with location and orientation sensors, partitions information access in a literally spatial manner, so that bringing the device in proximity to different parts of an answering machine, for example, might cause the palmtop to display the number, origin, and content of recently received messages [10].
- *ClearBoard*. The *ClearBoard* implementations (I and II) of Hiroshi Ishii and Minoru Kobayashi provide a minimal, transparent pane of mediation between two collaborators in disparate locations; *ClearBoard* uses video underlays in a drawing application (for example) in which each participant sketches directly on the workspace-pane while seeing the remote participant and her precisely-registered contributions to the evolving sketch 'just on the other side of the window' [11].
- *ALIVE*. Bruce Blumberg's *ALIVE* system is a focused implementation of the 'magic mirror' idea in which a large rear projection screen shows a live video image of the user himself (and the immediate physical surround) into which are composited real-time computer-generated elements, many of them autonomous animated creatures that react to and interact with the human participant [12].
- *ambientROOM*. Extending the *metaDESK*'s materialisation of interface a step further,

Hiroshi Ishii's *ambientROOM* invests a variety of physical objects and processes with direct yet poetically symbolic responsibility for output and input. In this office-like environment, network traffic can be represented aurally as the patter of rain while the evolving population of an external, communal office space is distilled into the delicate scintillations of an abstract pattern projected on a wall [9].

- *Wearable Computing.* Steve Mann and Thad Starner (among others) employ head-mounted displays and portable computational resources to move traditionally desk-bound systems into the perambulatory realm of immediate personal space. Experiments like Mann's 'live' image capture, processing, and dissemination and Starner's text-based 'Remembrance Agent' intend to provide extended access to various information sources and result in a kind of augmentation of the human informational memory [13, 14].
- *Personal Information Architecture.* Michael Hawley's Personal Information Architecture Group is building a model for reactive domestic and workspace environments, in which partially specialised interface schemes are built into everyday objects; while pervasive, small-scale networking and knowledge about the identity of other nearby system constituents and people makes possible (for example) FedEx packages that trigger the retrieval of shipping information when brought near a generalised display, and potted plants that – loitered near – whisper stock quotes, urgent phone messages, and so on to the loiterer [15].

The Luminous Room

A description of the *Luminous Room* begins most effectively with an inventory of its literal technical components and their general construction. With this in place we will then be able to discuss what is being done with the system, what will be done with the system, what could be done with the system, and what, in a philosophic sense, the idealised *Luminous Room* system is *for* and *about*.

System Construction

The physical apparatus of the *Luminous Room* consists of a high-resolution video projector mounted near the ceiling of an otherwise ordinary room in a gimbal comprising mechanically- and optically-rotational elements. The gimbal is driven by computer-controlled servo-motors, so that the projection can be aimed in any direction (including toward the ceiling itself). The projector is fed directly by a high-performance graphics computer, which among other things pre-warps its dynamic output to provide positionally independent, undistorted images despite the oblique projection angles often required of the system. Meanwhile, a miniature colour CCD camera is positioned coaxially with the projector's lens, and is subject to the same set of optical and mechanical manipulations as the projection itself. With the camera's lens also chosen to image the same solid angle as is projected, the system is in a perpetual state of 'looking where it's going' – moreover, this live video image of the region of the room currently addressed by the projected output is in fact the principal input to the system. Additional or duplicate input modalities, such as speech recognition and locked-off video cameras that image fixed parts of the space, are accessible to the system.

Before addressing particulars of the software components, we can already understand from the foregoing that the *Luminous Room* system is an instrument capable of painting dynamic information onto any surface in a room; further, it is implicitly capable of reactive dialogue with human users present in the space, inasmuch as the system not only paints but 'watches'. If we additionally assert that an understanding of both the geometrical and the *semantic* disposition of the containing room (including descriptions of non-architectural elements, like furniture, fixtures, various portable articles) is available to the system, then we have an apparatus capable of presenting information in a spatially relevant way. This means that the room's surfaces can be treated more meaningfully than as mere convenient projection surfaces: for example, a table should not be considered simply a region of the floor's expanse; information projected on the table will in most cases need to respect its edges as hard boundaries and not 'fall off' (especially as information should have been placed there because it needed to be *on the table*, not just because it needed to be somewhere). We will engage this idea in more detail later.

Current implementation

The *Luminous Room* system is at present in an early stage of development; the hardware described

above is by and large in place, and a core of software has been built.

Software structure. In particular, rather than agglomerating more and more functionality into a single monolithic program, the software partitions its work into a section that handles the machinations of controlling the gimbal to aim the projector, of adjusting the prewarping to accommodate the obliquity of the projection's intersection with each surface, of rendering portions of the image to be projected, of queuing and synchronising acoustic events, and of performing all manner of low-level tasks necessary to the literal mechanical and optical operation of the *Luminous Room*; and then any number of 'application' sections that provide whatever high-level functionality is desired and which treat the apparatus, as represented by the first software section, simply as a service to be exploited in the pursuit of their goals. Finally, a separate coordination section mediates among the single low-level module and multiple application modules, resolving possible overlaps and conflicts of requests for use of the basic system. These relationships are shown in Fig. 1.

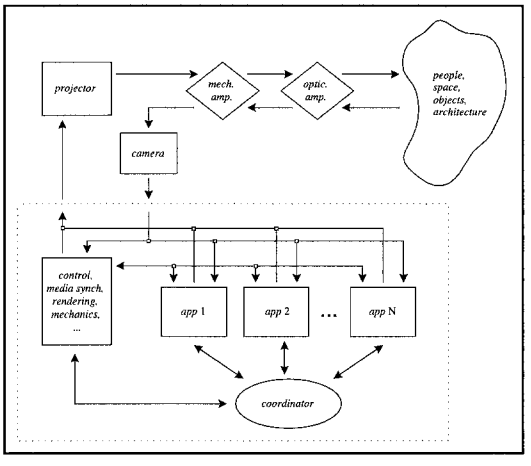

Fig. 1. Luminous Room system diagram.

Optomechanical design: A requirement of the *Luminous Room's* physical apparatus is that it be able to address the entire volume surrounding it– that is, walls, floor, and ceiling should all be possible sites of information interaction, with no 'blind spots' (projective or video-visual) anywhere on the sphere of its 4π steradian angular purview. A straightforward mechanical scheme accomplishing this would place the projector in a two-degree-of-freedom gimbal whose principal motion permits a 360° rotation about a vertical axis (providing longitudinal access) and whose subordinate motion permits 180° of latitudinal, 'top-to-bottom' rotation about a horizontal axis. Given the still substantial size and weight of current video projectors, the unavoidable prospect of the apparatus's sweeping out a full solid sphere during typical operation will be seen both as spatially prohibitive (and potentially dangerous to taller participants) and as inertially challenging for the on-board stepper- or servomotors responsible for such rotation.

Our design is a variation on this approach in which the secondary (latitudinal) rotation is accomplished optically rather than mechanically. Here, while the projector is still affixed to a stage that swivels 360° in the horizontal plane, the emerging image is caught by a mirror lying directly forward along the projector's optical axis but canted at 45° to it. This lightweight front-surface mirror can then be rotated along that same optical axis to sweep the image from directly overhead, through the horizon, to directly underneath. The small video camera that is the main input to the system is oriented to 'look into' this mirror as well, so that its field of view is always aligned with that of the projection. All this is sketched in Fig. 2 (as seen from below – the components shown actually depend from the rotational stage, which in turn depends from the longitudinal servomotor attached to the ceiling).

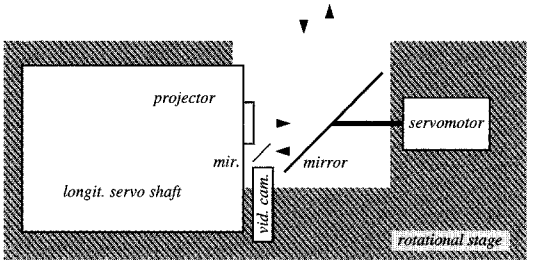

Fig. 2. Luminous Room optomechanics.

The operating volume required of this arrangement is significantly smaller than that of the fully mechanical gimbal and allows the rig to be placed much closer to the ceiling. A small incurred complexity is the gradual rotation of the resulting image about its center: while an image aimed by this system precisely toward the horizon is simply flipped left-to-right, adjusting the mirror to raise the image toward zenith or lower it toward nadir results not only in the desired gross vertical translation but also in a progressive local tipping of the image, ending with a ±90° rotation for positions fully above and below. At present we compensate

for this simply by pre-rotating the digital image fed to the projector by the calculably appropriate opposite-direction amount. (Note that the video stream acquired by the projection-aligned video camera will likewise be adversely rotated and will require similar digital treatment.) An alternate, non-computational means of achieving this corrective image pre-rotation is the interposition of a Dove prism between the projector and the mirror, with the prism ganged to the mirror to rotate in the opposite direction at half the mirror's angular velocity.

Applications: Two demonstration applications have so far been constructed. Input to each is performed literally by hand, the user simply roaming and gesturing along a wall of the room, and with no tethering, wiring, or augmentation of the user. Both applications run simultaneously.

In the first application, the user is able to demarcate a region of the room within the current bounds of projection and 'cut out' that region, creating a free-floating live image (the *live* character of the excised image is most evident if there is ongoing motion within the real region that it represents). Any number of pieces of the environment may be cut out in this way; the resulting images may be dragged about (again, purely through manual gesture) and rearranged, each continuing to present an active mirage of its physical location of origin. There are two further elements in this scenario: one is a real wastebasket, recognized and tracked by the system, that is an operative digital participant: images manoeuvred into it are terminally dispatched. The other is a large vase, also visually tracked by the system, which serves as a kind of 'genie bottle' – the user can pull a handful of live image chunks into the vase for storage, move the vase at will, and then cause it to disgorge its digital contents in some new location. (These manoeuvers are illustrated in Figs 3 and 4.)

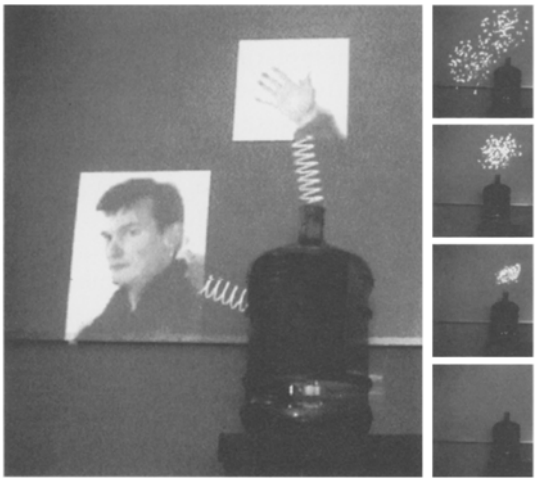

Fig. 3. Luminous Room optomechanics.

The prevalent feeling while exploring this scenario – a kind of real-world cut-and-paste – is a bit strange: the experience of ripping out duplicate pieces of animate reality and moving them around has a certain uncanny aspect to it.

The second application is a chess game, in which the system is able to recognise the presence of a physical chessboard on which the user has set up the black pieces. At the first appearance of the board, the system populates the appropriate

Fig. 4. The vase is moved to a new location; Its digital contents are then disgorged.

sixteen squares with its own pieces (Fig. 5.); these are of course white, but are virtual instead of corporeal. The game ensues, the user moving her physical pieces and the system responding with virtual, projected moves. As play progresses the board itself may also be moved about the space; the system's tracking keeps its pieces in their proper places.

Further scenarios

We are building a corpus of *reference access* applications: the appearance of a lambent recipe upon requesting 'Julia Child's Lemon-Basil Glaze,

Fig. 5. A physical chessboard is populated with the system's virtual pieces.

right here' when one's poultry-besmeared hands would otherwise stain a text is of clear value. A scrollable and searchable telephone directory that appears automatically near a picked-up phone makes sense. But just as important are a less-obvious group of applications that address *physical* information. 'Measure', indicating a slab of balsa wood, causes a ruler to be projected and edge-aligned; 'mark at seventeen inches' provides a bright cut-here line in the correct place, to be traced and sawed along.

There is a class of *connective* applications: these can function to bring together disparate spaces, as is the case with the 'alocational videophone' application, which allows participants in a sight-and-sound phone conversation to wander about a space; with the image of a distant conversant fluidly following the user about according to the user's automatically-tracked position, the 'stay-in-your-seat' stiffness that cripples normal videoconferencing is avoided and the potential for more natural, casual interaction provided. The *Real-World Cut-and-Paste* application already described is another example of this – it's easy to imagine an integrated 'Luminous House' in which an extended version of this application would allow a concerned father to 'cut out' the volume of an infant's room containing the baby-bearing crib and bring it along with him as he attends to other activities elsewhere in the house. There are also plans to demonstrate a second kind of connectivity afforded by the *Luminous Room* system: the bringing together of disparate *times*. With the appropriate high-density, random-access storage scheme it will be possible to provide a service of the sort that responds to commands like 'show me what happened here in this corner between four and five this afternoon'– video played back from the position of its recent-historical capture will reanimate that portion of the room, and if played considerably faster than real time will permit a rapid understanding of whatever activity may have brought the region to its current state.[1]

We also envision a category of *background-channel information* applications. In one such (which we have planned as 'Eloquent Wainscotting' and are implementing), the generally disregarded boundaries of a wall are invested with terse, unobtrusive summaries of larger information sources. Adopting a correspondence between the top, bottom, left, and right edges of the wall and (respectively) 'sky, earth, past, and present/future': weather information is displayed above, traffic below, news items to the left, and an appointment calendar to the right. The display of each kind of information is, again, subtle, thin and strip-like along its edge, and highly stylised. The upper-right corner is taken to represent the exact present, so that weather information nearest it along the top is associated with the immediate future while more distant predictions are arrayed progressively to the left, the whole scrolling imperceptibly to the right with the passing of time. Similarly, appointment representations are organised chronologically along the right edge, with impending obligations rendered nearest the top (and with additional attention-catching intensity); the entire strip likewise drifts upwards, toward 'now'. Finally, any of the edges may be 'drawn out' with a sweeping gesture that expands them to a foreground-attention, wall-sized display of full detail and complexity, in much the same way that pre-video-compositing weathermen yanked a succession of spring-rolled maps into the television frame from above.

Finally, *object-associated* applications are those which, like the chess-playing system, address the meanings and activities surrounding particular physical objects (or classes or collections of objects).

System Philosophy

As we conceive it, the *Luminous Room* embodies a certain philosophy that concerns *how* people

[1]The unequivocal ancestor of this application is Michael Naimark's splendid 'Displacements' installation, in which a rotating film camera captured a minute or so of activity in a populated and furnished room and then reprojected its having-been-developed optical contents onto the same room, now painted matte white. The effect is tremendous.

should be able to interact with computers. This philosophy in turn implies an idealised form for the *Luminous Room* system, one that is free from the constraints imposed by technological juvenility. In the system's current implementation, for example, shadows are a predictable but unavoidable shortcoming – a user can easily occlude the projection and extinguish the very information he had meant to examine. At the same time the projector's finite display angle means that information cannot be everywhere in the system's space at once. But the *Luminous Room* is not about high-resolution video projectors; a conjectural layer of transparent but addressably fluorescent material distributed pervasively about an environment would accomplish most of the same ends and sidestep many of the existing system's optical problems.

The *Luminous Room* is, instead, about untethered interaction – the possibility of accessing and manipulating information in any place that may be instantaneously convenient. It is also about pursuing this in particular places that may be specially appropriate, places where the information in question is conceptually already resident or resonant. The *Luminous Room* is about conducting digital dialogue very much on one's own spatial and physical terms.

These are the general tenets of the *Luminous Room's* conception. What follows are four notional roles played by the system.

LR as omnipresence of information

In its most obvious guise, the *Luminous Room* system enacts a thorough delocalization of information handling and display, both in the sense of *accessing* general information at any arbitrary position throughout a space and in the sense of *associating* relevant information with specific positions in a space. The first case is suggestive of Weiser's *Ubiquitous Computing*, while the second seems redolent of systems like Feiner's *KARMA* and Fitzmaurice's *Spatially Aware Palmtops*.

Ubiquitous Computing intends to achieve its ubiquity through a vigorous replication of hardware ("making many computers available throughout the physical environment" [16]) married to careful design that allows the technology to become universally reliable but fundamentally unobtrusive – that is, profitably taken for granted ("the most effective kind of technology ... essentially invisible to the user" [16]). As has been alluded to, one critical conceptual effect of this approach is that information becomes quantised into units of hardware: information is 'contained in' the system's *pads* and *tabs* which, though numerous, are individual portals through which communication is conducted. The *Luminous Room* proposes to furnish the same density of access – i.e. anywhere and everywhere the user may happen to be – but elects instead to make the environment itself the medium of interchange, so that information remains unquantised and is portallessly 'uncontained'.

Both Feiner and Fitzmaurice convincingly put forward the idea of an 'information field', an arrangement of spatial context that makes particular information uniquely available in the vicinity of the objects and places to which it intimately refers. Feiner's head-mounted, structurally-aligned means of access to this info-field has the more literal approach, overlaying annotations for the private apprehension of its user so that pieces of environmental geometry appear to be extended, or X-rayed, or temporally evolved. Fitzmaurice's system interprets the field symbolically, presenting locationally-triggered ancillary information, textual annotation, graphically related associations, and so on; but not in general attempting to spatially 'line up' its information with the underlying reality. In different modes the *Luminous Room* takes both these tacks: knowledge of the structure of the environment allows the system to augment regions pictorially as required; but it is just as often desirable to simply place relevant information in the proximity of its referent.

In both cases – location-independent access of general information and location-dependent presentation of an information-field – there is a common distinction between the approach taken by the *Luminous Room* system and that taken by its complementary practitioners, and it is the difference between 'looking at' and 'looking into'. This leads to one of the principal convictions guiding the design of the *Luminous Room*: that physical apparatus should not mediate or intervene, but that the user should look directly at information.

LR as digital participation for Ill-prepared objects

It is reasonable to imagine that no single ABC system will achieve sole eventual dominance, but rather that, due to the context-dependent applicability exhibited by each in turn, some collection of available systems and styles will in the long-run emerge as the future's computing environment. At the same time, it is fairly clear that the ability to make both sense and use of physical objects in real settings is a central motif in the landscape of ABC

systems. Nonetheless, and despite the projected best efforts of enhanced-object enthusiasts, it remains unlikely that any significant fraction of the world's human-manipulable mass will acquire the wires or tags or transceivers necessary to attain digital citizen status.

To be sure, those objects that are digitally enhanced, if well selected, will represent an important core of tomorrow's whole-ABC scenario; and, indeed, the more the better. The idealised *Luminous Room* system extends significantly the portion of the physical world accessible, embracing otherwise digitally mute objects. In the *Real-World Cut-and-Paste* application already implemented a trash can and a vase are recognised by the system and persuaded to reprise their normal corporal functions in the digital realm. There is, similarly, no reason why a computer should fail to exploit a perfectly good chessboard; here the existing *Luminous Room* makes use of an object in a curious way that lets an activity straddle the gap between real and virtual. The human and the computer address the abstract components of the undertaking in their respective domains, but are able to use the chessboard as a common instrument for expressing both output and input.

The most basic quantum of meaning that can be extracted from pieces of the real world depends on the recognition (or assignment) of *identity*, and so to the extent that any object can be passively identified by the *Luminous Room*, it can also be accorded an important kind of digital participation. Additional pieces of state, like position and orientation, when discernible only increase this usable meaning – the enthalpy of the information – geometrically. The *Luminous Room* system is able to accomplish this without requiring augmentation of its objects. Of course, it should be stated that the system can make equally profitable use of objects that *are* digitally enhanced and whose identity and state have been relayed by whatever external systems are responsible for them; it does not ultimately matter through what channel – remotely- or self-identified, internally or externally provided – such information comes.

LR as universal adaptor

Beyond its principal function as a medium for interaction and information manipulation, the *Luminous Room* system can also behave as a global connector, an 'information impedance matcher' that externalises translation (and offers control over that translation) throughout the universe of ABC systems, non-ABC systems, and human users.

As an example, we have noted already that in providing an intimate framework for highly personalised information access *Wearable* systems are necessarily exclusive; it makes no real sense for a *Wearable* user to say 'look at this', particularly to someone who is not also draped with a similar system. However, in the presence of a *Luminous Room* system the intent of such a statement can readily be carried out, the image or information in question being re-presented externally and to any number of viewers. The particular manner in which transactions of this sort are designed will have a great deal of bearing on their perceived meaning and import: should a replica of the *Wearable* user's head-mounted view simply be replicated on the nearest wall without further annotation, should some symbolic depiction of the flow of information 'out of' the *Wearable* user and onto the floor precede or accompany the public *display* of information on the floor, should the information appear on the *Wearable* user herself?

In this case the *Luminous Room* can be seen acting as a connection between some other system and one or more humans; there may also be value in allowing the Luminous Room to visualize the flow of information between two systems, generating an external representation of transfer or coupling processes. In an environment in which several concurrent ABC systems provide an overlap of services and resources, it may be occasionally necessary or desirable for a human to 'direct traffic' – and so we imagine the system providing not only a visual portrayal of the flow but also the means to control, direct, monitor, divert, reconnect. The *Luminous Room* is seen here as a distributed 'conductor's podium' from which the operations of an environment full of other information systems can be managed and orchestrated as necessary – an omnipresent control panel. If we consider the same functionality from a slightly different view, the *Luminous Room* provides an ideal interface *service* to which *other* information systems can subscribe, relying on it to help them extend necessary controls to a user, and trusting it with organisation, unification, and arbitration of the predictable multiplicity of such subscriptions.

LR as global referencer

A basis for human communication (in- and external) is a mind's ability to refer to things outside itself. The simplest kind of externalised communication can proceed through literal pointing, a procedure for designating that works across and despite a separating distance. Verbal language

supplies the further power of progressively abstract reference: nouns that can refer to things that aren't immediately present; verbs that can refer to things that aren't things but rather actions or processes; pronouns that can refer indirectly to many different things by referring directly to other nearby pieces of language; chunks of speech that can refer to things that aren't ever pointable-to, because they are ideas and thus incorporeal; and so on.

As a system which incorporates a knowledge of its own physical environment and a facility for perceptibly modifying that environment, the *Luminous Room* is possessed of a unique expressiveness. At the most primitive level, the system can 'fetch and present' pieces of the world (in the manner illustrated by the *Real-World Cut-and-Paste* application), offering these visually as parts of some compound 'statement'; this is rather like the method resorted to by the cinematic dog whose inarticulateness-thwarted intelligent communication must instead be carried out through a trot home with an edifying chunk of the particular well down which Timmy has fallen this week. Alternately, the system can indicate artifacts around it in a variety of gestural ways: there is, for example, little more unmistakable and indisputable than an illuminated arrow pointing at some object in question, if absolute bluntness is what's required ('the keys you've misplaced are right here').

Beyond that, the *Luminous Room* can act to perform visual annotation of parts of its surroundings – placing carpentry-assist marks on lumber, didactically extending some wall-hung Bosch to show the entire *Garden of Earthly Delights* from which it was cropped, animating the twisting trajectory of a piano that must be moved through a challenging doorway. Such deeds can be seen as an equivalent to the construction of sentences in which 'tangible' nouns figure.

The Question of Architecture

At one time the structure of a house was concerned quite directly with the simple aim of discouraging the elements' advance and providing for its denizens a measure of functional comfort that surpassed what they could expect from the great outdoors. To be sure, a house may also serve to collect and delimit a unit of family, to stand as a marker of property without and container of property within, to establish through its size or allotment of grandeur some position within a social hierarchy; these are its symbolic affairs. But if we momentarily restrict our understanding of *house* to the set of functions that accommodate the tangible needs and activities of those within, then we can point to a suggestive progression of enhancements – a chronological parade of infrastructural improvements hinting at a hypothetical destiny for the things that ABC would offer.

Chamberpots and free-standing claw-footed tubs that waited to be filled by bucket-brigade were obviated by pipes charged with the delivery and elimination of water. The precarious canopy of electrical distribution wires in houses built without foreknowledge of an imminent world of convenience gave rapid way to discreet channels, planned for and tucked out of sight at construction-time, through which currents could complete their appointed rounds. Electric fans for summer and heating registers for winter would be eliminated in favour of internal ducts that exhale the hot or cold breath of machines exiled to the basement. We now build with distribution of telephone and cable television signals very much in mind, and these modest information tubes will soon enough be joined in every dwelling that rises by digital network conduits. So it is that a new house can expect to be given a circulatory system, a respiratory system, a warm-blooded thermal regulation system, and a nervous system. The right collection of ABC implementations might expect to provide a brain.

At the same time, each of these additions has asked very little of the fundamental structure of the house: a little wall area here, a few extra holes there, a bit of clearance between floor and ceiling of this room and that. A well-integrated set of ABC installations would require little more. Our houses have digested these technologies and capabilities quite easily; no especially recognisable mutation or evolution of the underlying architecture has been necessary.

The *Luminous Room*, meanwhile, is information in architecture-soluble form. It is, naturally, also architecture-*digestible* because, in its ideal manifestation, whether mounted on the ceiling as a small optical nozzle or coated universally on walls and floors as an addressable luminous layer, its actual apparatus would be less evident or invasive even than a sink, telephone, or thermostat. But it is architecture-*soluble* because, unlike the plumbing and wiring and ducting whose successor it would be, it is specifically about the space in which it is resident. In short, it is the first system that addresses architecture not so much as a topic (although some *Luminous Room* applications

do take this as their subject) but specifically as a medium of active interchange, a forum for active communication.

It is more or less axiomatic that any architectural project is executed in service of some particular function: stadia and concert halls accommodate people and the events they desire to see en masse; grain silos and water towers contain and transport aggregate inanimates; factories and Bauhauses surround and inform processes of production. In every case, what is to happen within dictates most elements of gross form. And yet it might be argued that almost never is the *experience* of architecture itself the primary *function* of an architectural construction. By *experience of architecture* we mean specifically *the set of transactions – whether passive or active – that occur at an architecture's surfaces.*

The reason that the experience of architecture and its primary function are generally disjoint (and this brings us at last to the central point) is that what is *built* is positive space, but what is *used* – the place where all activities happen – is negative space.[2] Obviously, architectural process takes negative space very much into account at the time of its more explicit design of positive space; but the fact remains that humans have so far been restricted for active use to the holes in the positive space, to what the positive space bounds and surrounds and leaves behind: the sofas, stoves, desks, Lego, beds, pianos, and computers must all be placed in the negative space. Of course, we do have access to the positive space, but in a very constrained way that leads generally to passive ornamental or decorative use: frescoes, rose windows, mass-reproduced Impressionist paintings, and (non-decoratively and slightly more actively) post-it notes are all viable positive space implements. At the extreme, Robert Venturi has asserted that in Las Vegas (and, by implication, Manhattan's Times Square, Tokyo's Ginza and Shinjuku) the enormous, self-luminous signs that precede or are attached to the buildings to which they usually refer are not only a legitimate component of the region's architecture, but are in fact the dominant and defining characteristic of such architecture [17]. And indeed, to the extent that such architecture is experienced from the point of view of a pedestrian or motorist (who is in the infinite negative space of 'outside') here the experience *is* the function. Still, this is a passive experience – 'watching only'.

What the *Luminous Room* and similar systems provide, through their architecture-solubility, is an active engagement with positive architectural space. Use of the Luminous Room system entails approaching walls, pointing at ceilings, sitting on floors, and attending to them all: these are the places where information appears, where it is available for manipulation. While this direct involvement with structure, with the 'primary architectural space', is not itself a goal of the *Luminous Room's* design, it is a critical consequence of its operation.

Is it possible, then, to foresee a new kind of architecture that responds at a much more basic level to information flow? As our existences are given over more and more to the pursuit of digital activities – more, certainly, and in a more significant way than they are given to plumbing-mediated undertakings, for example – might not that predominant pursuit be overtly supported by the surrounding architectures? If, in particular, we find ourselves reliably served by ABC information systems like the *Luminous Room* that conduct their business as respectful intimates of architecture, then it seems quite conceivable that design practice could in turn come to inflect architecture back toward that intimacy, in support of an even more fluid use of information by humans and of architectural space by information. It is far too early even to speculate sensibly about what forms might result from such a symbiosis. But we are now seeing the first emergence of information systems that make a strong effort to find a place in the physical world, and among them are projects like the *Luminous Room* that specifically occupy the architectural surfaces which bound much of human activity; a responsive architecture will inevitably come to acknowledge this.

Conclusion

Research in systems that deliberately set aside the unimaginative status quo of human–computer interaction has, through the acknowledgment of shared objectives and philosophies, begun to coalesce into a broad but discernible field, here called *Antisedentary Beigeless Computing.* A new

[2]Cathedrals are a close miss: the mood of overwhelming and awe-injecting architecture will indeed critically inform a religious proceeding. Then again, at the end of the day the cathedral still contains (and so is about containing) the religious proceeding. Then again again, it could easily be argued that a magnificent place of worship is as much a part of its associated god's conception (and so a part of that god's worship) as anything else – and that it therefore is about its own experience.

addition to the stable of this work, the *Luminous Room* project, occupies one extreme of the ABC envelope in which information access is both fully distributed throughout a surrounding architectural volume and is accomplished entirely without the direct manipulation of specialized hardware.

The *Luminous Room* project should ultimately be taken not so much as a system of opto-electro-mechanical projective hardware and coordination-and-control software as a collection of applications delineating a philosophy of what can and ought (and also ought not) to be provided to humans given the existence and prevalence of such hardware and software facilities. It is the particular province of the *Luminous Room* system to provide for omnipresent information access and manipulation throughout a (possibly quite extended) space; to make physical objects of all sorts, whether explicitly digitally readied or not, integral subjects and objects of this manipulation; to mediate among other ABC systems whose otherwise invisible transactions may benefit from 'literal' manifestation and human control access or that may require the services of its heavily display-oriented capacities; and to do all of this with explicit knowledge of, integration with, and reference to a containing architectural space.

There is a very natural connection between the collective concerns of ABC research and those of the field of architecture; the overlap has so far been little addressed or explored. But as development of ABC systems matures and passes beyond the merely illustrative or experimental, that link will necessarily become explicit and an applicable understanding of it elaborated. At the moment, we expect that systems like the *ambientROOM*, PIA, and the *Luminous Room* (all of which depend on viable architectural integration and so cannot ignore or gloss over this link) will through their own growth be the first bodies of work also to offer theories of how information and architecture can most usefully be amalgamated.

References

1. Underkoffler J. Antisedentary Beigeless Computing. Personal Technologies. Springer-Verlag, 1997. 28–40
2. Bolt R. The human interface: Where people and computers Meet. Lifetime Learning Publications, Belmont, 1984.
3. Bolt R. 'Put-that-there': Voice and gesture at the graphics interface. In: Proceedings of SIGGRAPH, 1980. 262–270.
4. Bolt R. Gaze-orchestrated dynamic windows. In: Proceedings of SIGGRAPH, 1981. 109–119.
5. Krueger M. Artificial reality II. Addison-Wesley, Reading, 1991.
6. Wellner P. Interacting with paper on the digital desk. Communications of the ACM, July 1993; 36(7), 87–96.
7. Feiner S, MacIntyre B, Seligmann D. Knowledge-based augmented reality. Communications of the ACM, July 1993; 36(7).
8. Weiser M. The computer for the 21st Century. Scientific American, September 1991; 94–104.
9. Ishii H, Ullmer B. Tangible bits: Towards seamless couplings of people, bits, and atoms. To Appear In: Proceedings of CHI '97.
10. Fitzmaurice G. Situated information spaces and spatially aware palmtop computers. Communications of the ACM, July 1993; 36(7).
11. Ishii H, Kobayashi M, Arita K. Iterative design of seamless collaboration media. Communications of the ACM, August 1994; 37(8).
12. Maes P, Darrell T, Blumberg B, *et al.* The *ALIVE* system: Wireless, full-body interaction with autonomous agents. ACM Special Issue on Multimedia and Multisensory Virtual Worlds, Spring 1996.
13. Starner T, Mann S, Rhodes B, *et al.* Wearable Computing and augmented reality. Presence, MIT Press, Cambridge, 1997.
14. Mann S. Wearable computing: A first step towards personal imaging. IEEE Computer, February 1997.
15. Hawley M, Poor R, Tuteja M. Things that think. Personal Technologies, Springer-Verlag, 1997. 1, 13–20
16. Weiser M. Some computer science problems in ubiquitous computing. Communications of the ACM, July 1993; 36(7).
17. Venturi R, Brown DS, Izenour S. Learning from Las Vegas: The forgotten symbolism of architectural form. MIT Press, Cambridge, Massachusetts 1977.

Correspondence and offprint requests to: John Underkoffler (jh@media.mit.edu) MIT Media Laboratory, 20 Ames Street, Cambridge, MA 02139, USA.