

Transcending the Individual Human Mind

Creating Shared Understanding through Collaborative Design

Ernesto Arias^{1,3}, Hal Eden^{1,2}, Gerhard Fischer^{1,2},
Andrew Gorman^{1,2}, and Eric Scharff^{1,2}

¹Center for LifeLong Learning & Design and Institute of Cognitive Science

²Department of Computer Science, ³College of Architecture and Planning
University of Colorado, Boulder

{ernie, haleden, gerhard, agorman, scharffe}@cs.colorado.edu

Abstract

Complex design problems require more knowledge than any single person possesses because the knowledge relevant to a problem is usually distributed among stakeholders. Bringing different and often controversial points of view together to create a shared understanding among these stakeholders can lead to new insights, new ideas, and new artifacts. New media that allow owners of problems to contribute to framing and resolving complex design problems can extend the power of the individual human mind.

Based on our past work and study of other approaches, systems, and collaborative and participatory processes, this paper identifies challenges we see as the limiting factors for future collaborative human-computer systems. The Envisionment and Discovery Collaboratory (EDC) is introduced as an integrated physical and computational environment addressing some of these challenges.

The vision behind the EDC shifts future development away from the computer as the focal point, toward an emphasis that tries to improve our understanding of the human, social, and cultural system that creates the context for use. This work is based on new conceptual principles that include creating shared understanding among various stakeholders, contextualizing information to the task at hand, and creating objects-to-think-with in collaborative design activities. Although the EDC framework is applicable to different domains; our initial effort has focused on the domain of urban planning (specifically transportation planning) and community development.

Content Indicators

Categories and Subject Descriptors

D.2.2 Software, Tools and Techniques; H. Information Systems; H.1.2 Models and Principles, User/Machine Systems, H.3 Information Storage and Retrieval; H.5 Information Interfaces and Presentation; H.5.1 Multimedia Information Systems; H.5.2 User Interfaces, H.5.3 Group and Organizational Interfaces; I.3.6 Computing Methodologies, Methodology and Techniques; J.4 Computer Applications, Social and Behavioral Sciences

General Terms

Human Factors, Design

Keywords:

collaborative design and knowledge construction; symmetry of ignorance; integration of action and reflection spaces; integration of physical and computational environments; open systems; design support systems; distributed cognition

Table of Contents

| | |
|--|----|
| <i>Abstract</i> | 1 |
| <i>Content Indicators</i> | 2 |
| Categories and Subject Descriptors: | 2 |
| General Terms | 2 |
| Keywords: | 2 |
| <i>Introduction</i> | 4 |
| <i>Challenging Problems for the Future of Human-Computer Interaction</i> | 4 |
| <i>The Envisionment and Discovery Collaboratory (EDC)</i> | 8 |
| A Scenario: Creating Shared Understanding through Collaborative Design. | 10 |
| The Conceptual Principles behind the EDC..... | 14 |
| The EDC Architecture..... | 15 |
| The Integration of Action and Reflection. | 15 |
| The EDC as an Open System..... | 16 |
| <i>Assessment</i> | 17 |
| <i>Future Work</i> | 18 |
| <i>Conclusions</i> | 18 |
| <i>Acknowledgments</i> | 19 |
| <i>References</i> | 19 |

Table of Figures

| | |
|---|----|
| Figure 1: The EDC as a convergence of systems | 9 |
| Figure 2: The Current Prototype of the EDC | 9 |
| Figure 3: Placing Objects in the Action Space..... | 10 |
| Figure 4: Drawing Roads | 10 |
| Figure 5: Specifying Travel Preferences..... | 11 |
| Figure 6: Retrieving Constructions from the Reflection Space..... | 11 |
| Figure 7: The Color of the Bus and X's Surrounding the Bus Stop Enable Participants to Visualize Information in the Simulation | 12 |
| Figure 8: A Web-Based Survey Allow Participants to Parameterize the Bus Stop's Attributes..... | 13 |
| Figure 9: The Reflection Space Combines Domain Expert Knowledge with an Extensible Discussion Forum to Provide a Mechanism for Learning on Demand and the Construction of New Knowledge .. | 14 |
| Figure 10: The EDC Layered Architecture | 15 |
| Figure 11: Blending Action and Reflection..... | 16 |

Introduction

Human-computer interaction (HCI) research over the last twenty years has made fundamental contributions to the creation of new paradigms and new forms of working, learning, and collaborating in the information age. Its major emphasis has been to develop new technologies (e.g., at the hardware, basic software and application levels), new interaction techniques (e.g., graphical user interfaces), new design approaches (e.g., user-centered, human-centered, work-oriented, and learner-centered design). Much of this research has emphasized and pioneered socio-technical approaches. In the process, HCI work has progressed from early concerns with low-level computer issues to a focus on people's tasks [Myers, 1998; Newell & Card, 1985; Norman, 1990]. The greatest progress in HCI research has been made at the operator and task level, where events are studied in time scales ranging from microseconds to minutes (and in some cases hours or days). At these time scales, the relevant theory is drawn from psychology and bounded rationality. The users considered were novices rather than skilled domain workers, which allowed researchers to do meaningful usability studies in the laboratory with undergraduates. As we enter the new millennium, we, along with others, claim that the major challenges of HCI will be at the design, system, technology, and media levels, where actions and changes take months, years, and decades. In long-term use settings such as these, the relevant theory to be taken into account will be primarily grounded in social and organization themes [Hutchins, 1994; Thomas & Kellogg, 1989].

In this paper, we first identify a set of challenging problems for HCI in the next millennium. We then describe our approach to address these challenges by focusing on the Envisionment and Discovery Collaboratory (EDC). A scenario is introduced that illustrates the current features of the EDC in a specific problem context, which grounds the discussion of the conceptual framework, the architecture, and the specific substrates of the EDC. We then briefly describe how our interaction with user communities has guided us in the assessment and iterative design of the EDC. We conclude by articulating some of the many remaining challenges of this approach for HCI in the future.

Challenging Problems for the Future of Human-Computer Interaction

Transcending the Individual Human Mind. The power of the unaided, individual mind is highly overrated: the Renaissance scholar no longer exists. Although creative individuals are often thought of as working in isolation, the role of interaction and collaboration with other individuals is critical [Engelbart, 1995]. Creative activity grows out of the relationship between an individual and the world of his or her work, and from the ties between an individual and other human beings. The predominant activity in designing complex systems is that participants teach and instruct each other [Greenbaum & Kyng, 1991]. Because complex problems require more knowledge than any single person possesses, it is necessary that all involved stakeholders participate, communicate, and collaborate with each other. For example, domain experts understand the domain concepts and practice whereas system designers know the technology. Communication breakdowns are often experienced because stakeholders belonging to different cultures use different norms, symbols, and representations [Snow, 1993].

In designing artifacts, designers rely on the expertise of others [Galegher et al., 1990; Resnick et al., 1991] by referring to textbooks, standards, legal constraints, and especially previous design efforts. Project complexity forces large and heterogeneous groups to work together on projects over long periods of time. Knowledge bases to support design should include not only knowledge about the design process but also knowledge about the artifacts of that process—parts used in designing artifacts, subassemblies previously created by other design efforts, and rationale for previous design decisions [Fischer et al., 1992]. Designers generally have a limited awareness and understanding of how the work of other designers within the project—or in similar projects—is relevant to their own part of the design task. The large and growing discrepancy between the amount of such relevant knowledge and the amount any one designer can possibly remember imposes a limit on progress in design. Overcoming this limit is a central challenge for developers of systems that support collaborative design [Nakakoji et al., 1998].

Exploiting the Symmetry of Ignorance. When a domain reaches a point at which the knowledge for skillful professional practice cannot be acquired in a decade, specialization increases, collaboration becomes a necessity, and practitioners make increasing use of reference aids, such as printed and computational media supporting distributed cognition. Design [Simon, 1996] is a prime example of such a domain. Complexity in design arises from the need to synthesize different perspectives of a problem,

manage large amounts of information relevant to a design task, and understand the design decisions that have determined the long-term evolution of a designed artifact. Design problems are wicked (i.e. ill defined and ill structured [Rittel & Webber, 1984]); they are moving targets that have resolutions rather than solutions [Arias & Schneider, 1999]; and the context in which these problems exist is by nature characterized by change, conflict, and multiple stakeholders [Arias, 1995]. In many cases, consensus is not achievable, and the best we can strive for is informed compromises emerging from *the symmetry of ignorance* [Rittel, 1984]—different aspects of knowledge crucial to the resolution of the problem carried in the minds of individual stakeholders as tacit knowledge. For example, this symmetry might represent different descriptions of the world or reasons behind conflicting arguments and goals among differing agendas in complex design problems.

Rather than viewing the symmetry of ignorance as an obstacle during design, we view it as an opportunity for the creation of new knowledge and new ideas (as observed by C.P. Snow: “The clashing point of two subjects, two disciplines, two cultures ought to produce creative chaos.” [Snow, 1993]). Having different viewpoints helps one discover alternatives and can help uncover tacit aspects of problems.

Exploiting the symmetry of ignorance requires putting owners of problems in charge [Fischer, 1994b], which will promote direct and meaningful interaction that involves people in decisions that affect them [Arias, 1996]. In order to bring important perspectives to the process of design, all stakeholders in the process should be designers and co-developers, not just consumers [Fischer, 1998]. End-users, as owners of problems, bring perspectives to collaborative design activities that are of special importance for framing problems. The existence of the symmetry of ignorance requires creating spaces and places that serve as *boundary objects* (shared objects to *talk about* and to *think with*) where different cultures can meet and collaborate. Boundary objects serve as externalizations [Bruner, 1996] that capture distinct domains of human knowledge. They have the potential to lead to an increase in socially shared cognition and practice [Resnick et al., 1991].

Accepting that most design problems are characterized by the existence of the symmetry of ignorance leads to a different view of expertise and learning. In these contexts, relevant knowledge, which needs to be drawn out of and synthesized from the perspectives and expertise of the contributors, does not already exist and cannot simply be passed on by those who have it to those who need it. Therefore, approaches are required that view learning as collaborative knowledge construction [Scardamalia & Bereiter, 1994] and expertise as a relative concept [Fischer, 1993]. This view is in sharp contrast to the teaching cultures of our schools [Illich, 1971], by which teaching is often “fitted into a mold in which a single, presumably omniscient teacher explicitly tells or shows presumably unknowing learners something they presumably know nothing about” [Bruner, 1996; Roggoff et al., 1998]. Likewise, the view of the domain expert as the sole source of design knowledge fails to recognize the fact that all stakeholders have important contributions to make.

Recognizing the Need for Externalizations in Collaborative Design. Distributed cognition [Norman, 1993] emphasizes that the heart of intelligent human performance is not the individual human mind in isolation but the interaction of the mind with tools and artifacts as well as groups of minds in interaction with each other. It is important to understand the fundamental difference between these two forms of distributed cognition. When distributed cognition is at work between the individual human mind and artifacts, such as memory systems, it often functions well because the knowledge an individual needs is distributed between her/his head and the world (e.g., an address book, a system of e-mail message folders, or a file system). But in the case of distributed cognition in operation among groups of minds, a group has no head, no place for the information about this distribution of knowledge to be available to all members implicitly—therefore externalizations are critically more important for collaborative design. Externalizations (1) create a record of our mental efforts, one that is “outside us” rather than vaguely in memory and (2) represent artifacts that can talk back to us [Schön, 1992] and form the basis for critique and negotiation.

A challenge is to integrate the various perspectives emerging from the symmetry of ignorance among articulate stakeholders. By supporting the process of reflection within a shared context defined by the task at hand, opportunities can emerge for enhancing the creation of shared understanding. This process melds the information that is collaboratively constructed into the problem-solving context, informing the process as well as the stakeholders and allowing them to participate from a more enriched and meaningful perspective [Brown et al., 1994]. It also enhances the quality of the designed artifact due to the synergy of interaction that draws out ideas and perspectives in a conversational manner. The resulting,

richly contextualized information is available for future stakeholders [Fischer et al., 1992] to draw upon, informing them not only about the surface level of the design, but about the deeper characteristics behind the design [Moran & Carroll, 1996].

Externalizations are used to extend our cognitive abilities [Engelbart, 1995; Norman, 1993] by allowing all stakeholders to engage in a “conversation with the materials” [Schön, 1983]. Our research has demonstrated that these “conversations” are very different in physical versus computational environments [Arias et al., 1997]. There is a growing interest in blending real-world artifacts with computational media [Eisenberg & Makay, 1996; Ishii & Kobayashi, 1992; Ishii & Ullmer, 1997]. Frequently, the design of interactive systems focuses *exclusively* on the capabilities provided by the dynamic nature of computational media. Yet physical models provide certain strengths not found in computational models. Rather than viewing this as a dichotomy—where one must choose between one or the other—HCI needs to explore the creation of combined physical and computational environments that use the strengths of each to augment the weaknesses of the other [Arias et al., 1997].

Contextualizing Information. If new HCI approaches, techniques, and systems are to be helpful, they should not be focused on producing more decontextualized information—most humans already have enough to occupy them from dawn to dusk. Rather, the emphasis should be on developments that take into account that human attention is the scarce resource [Simon, 1996] and help people attend to the information that is the most relevant for their task at hand.

In most situations, humans want to act—they do not want to study large information spaces (e.g., help information, design rationale) in the abstract [Moran & Carroll, 1996]. As they act, however, they experience breakdowns [Fischer, 1994c]. This leads them to reflect upon their activities, and in this context they explore information spaces associated with the activity. Schön calls this approach “reflection-in-action” [Schön, 1983], and in our own previous work we call it “making argumentation serve design” [Fischer et al., 1996]. This notion as well as our efforts to integrate action and reflection with critics and specification components in domain-oriented design environments have set our approach apart and have provided us with a unique foundation to create systems that “say the ‘right’ thing at the ‘right’ time in the ‘right’ way” [Fischer, 1994a] and support learning on demand [Fischer, 1991].

Supporting New Forms of Civic Discourse: From Access to Informed Participation. Another fundamental challenge for HCI in the next millennium is to invent and design a culture in which humans can express themselves and engage in personally meaningful activities. However, a large number of the new media are designed to see humans as consumers only [Fischer, 1998]. A prominent example of a consumer perspective was articulated by the director of research for Time Warner Entertainment in his closing plenary address at CHI ‘95. He challenged the HCI community with the task of designing a remote control to browse and efficiently select 500 or more TV channels. Solving this problem is of great commercial interest to industries that regard humans as the ultimate consumers—but is it a focal issue for HCI?

This emphasis on people as consumers is perpetuated in other perceptions of the future as well. The President’s Information Technology Advisory Committee’s report includes the call that “*The Nation must ensure that access to the benefits of the information infrastructure are available to everyone in our Nation*” [PITAC, 1999, p. 10]. While the universality of this vision is important, our claim is that more than *just* access is needed. An example of this broader vision was set forth by the President’s Council on Sustainable Development:

How can more than 261 million individual Americans define and reconcile their needs and aspirations with community values and the needs of the future? Our most important finding is the potential power of and growing desire for decision processes that promote direct and meaningful interaction involving people in decisions that affect them. Americans want to take control of their lives. [PCSD, 1996, p. 7]

The Council substantiates an increasing trend toward grass-roots, bottom-up efforts to address the impacts of growth (or decline) on the quality of life in U.S. communities. The nature and intensity of these impacts require difficult decisions on how to sustainably manage such growth in the future.

The broad challenge, then, is to move toward new forms of *citizen participation*. Certainly this challenge is not without its difficulties. For example, some of these include (a) the paradox that citizens cannot really be informed unless they participate, yet they cannot really participate unless they are informed [Brown et al., 1994]; and (b) that participation has limits that are contingent on the nature of each citizen’s situation, the issues, the problems, and the institutional designs [Arias, 1989], as well as the available technology

and media. One of the benefits of addressing these challenges is that informed participation leads to *ownership* and a stronger sense of community.

The challenge to the HCI community is to move beyond an emphasis on interaction that is solely focused on access to information to one that supports *informed participation*. This rests on the premise that one of the major roles for computational media is not merely to deliver existing and predigested information to individuals but to provide the opportunity and resources for design activities embedded in social debates and discussions in which all people can act as designers if they choose to do so rather than being confined to consumer roles.

Moving beyond Closed Systems. If HCI systems are to effectively support collaborative design, they must adequately address not only the problem situations, but also the collaborative activity surrounding the problem. By addressing real-world problems that are inherently ill structured and ill defined, the system must cope with problem contexts that change over time. In addition to the fluid nature of the problems themselves, the very process of collaboration among stakeholders further increases the ever-changing problem context. Because the issues that arise in these problems will depend on the background, motivation, and agendas of the participants, the problem will take different forms, depending on the collaborators. Designing systems to support the constantly evolving problem context as the collaborators work to understand, frame, and address it is an important challenge. Providing *closed systems*, in which the essential functionality is fixed when the system is designed, is inadequate for coping with such dynamic problem contexts. Creating a system with constrained functionality requires making assumptions about use that cannot be fully anticipated when the system is designed, because many of the issues come out only when a system is used.

Providing *open systems* is an essential part of supporting collaborative design. An open system provides opportunities for significant changes to the system at all levels of complexity. Enhancement and evolution of the system are “first-class design activities.” By creating the opportunities to shape the systems, the owners of the problems can be involved in the formulation and evolution of those problems through the system. The challenge for these open systems is to provide opportunities for extension and modification that are appropriate for the people who need to make changes. This is based on the following principles:

- *Software systems must evolve; they cannot be completely designed prior to use.* System developers cannot anticipate and design for every possible situation. Although it may not be possible to design “complete” systems, this does not mean that *all* aspects of a system must be constructed through user-directed evolution. In such a system, users would be unlikely to wish to spend considerable effort constructing even the simplest situations. Instead, designers must provide a *seed* for the system. The seed has an initial core functionality that can be readily applied to some situations and facilitates the construction of new situations. The seed must be designed to *evolve* over time, allowing users to make incremental changes to the core functionality when necessary. Eventually, designers and users may *re-seed* the system by incorporating pieces that were created during the system’s evolution into the core of subsequent systems. We have discussed this process model for evolution in greater detail previously [Fischer & Scharff, 1998b].
- *Systems must evolve at the hands of the users.* Giving the *owners* of problems the ability to change systems as they explore their problem leverages the insight into problems that uniquely belongs to those experiencing the problems. Many systems have explored the notion of end-user programming [Nardi, 1993], often focusing on providing mechanisms for nonprogrammers to change systems. Our focus is on end-user modification, where programming is just one form of modification necessary to evolve systems. The ability to specify goals and structure information are examples of other important modification tasks. Furthermore, the notion of “end user” need not be limited to someone who is not a programmer. Instead, it is important to provide different avenues for modification that are appropriate for different kinds of stakeholders.
- *Systems must be designed for evolution.* Extending an application in an initially closed design may be difficult because of the assumptions implicit in a system designed without extension in mind [Girgensohn, 1992]. A closed system with some extension capabilities will likely restrict what can and can’t change. Designing a system for evolution from the ground up, however, can provide a context in which change is expected and can take place. But because it is not known in advance what way a system will evolve, even the underlying assumptions behind an evolvable system may be

suspect. Therefore, it is important to design with an understanding of the nature of potential extensions, for some changes will always be more difficult than others.

- *Evolution of systems must take place in a distributed manner.* Systems must acknowledge the fact that users will be distributed both in space and in time. Distributed systems provide a framework for evolution in which all participants have the chance to contribute in a manner appropriate to their ability. The success of distributed open systems (as measured by their creation and continual growth by communities of users who are not obliged to extend the systems) is a testament to the efficacy of the distributed approach [Raymond, 1998]. The Educational Object Economy (<http://www.eoe.org>) and Gamelan (<http://www.gamelan.com>), repositories of resources for Java developers, are examples of systems that have grown largely through the participation of the community of developers [Fischer & Scharff, 1998a].

Understanding Motivation and Rewards. Computational support mechanisms are necessary prerequisites, but not sufficient conditions to motivate people to become part of a “design culture.” People must be motivated and rewarded for investing time and effort to become knowledgeable enough to act as designers. These rewards may range from feeling in control (i.e., independent from “high-tech scribes”), being able to solve or contribute to the solution of a problem, fulfillment of a passion to master a tool in greater depth, making an ego-satisfying contribution to a group, and/or contributing good citizenship to a community [Grudin, 1994].

Summary of Challenging Problems for the Future of Human-Computer Interaction. We have identified seven challenges that should be integrated into future HCI agendas. These agendas need to include the development of innovative information technologies to support collaborative design and learning in domains characterized by complex problems—in particular, they should include a basis for understanding *how* and *why* to:

- support distributed cognition in order to transcend the individual human mind,
- exploit the symmetry of ignorance by construction shared understanding,
- utilize externalizations to extend our cognitive abilities,
- contextualize information to avoid information overload and to increase opportunities for learning on demand,
- introduce and support the notion of informed participation because access, although necessary, is not sufficient,
- move beyond closed systems to support open, evolving contexts of complex design problems, and
- understanding motivation and rewards necessary to engage people in a design culture.

These challenges shift future development away from the computer as the focal point toward efforts that improve our understanding of the human, social, and cultural systems that create the context for use [Greenbaum & Kyng, 1991]. This vision and its conceptual understanding have guided us in the development behind the Envisionment and Discovery Collaboratory, an integrated environment for learning and design in which users discover and frame problems and construct new visions.

The Envisionment and Discovery Collaboratory (EDC)

To create a context for our study of shared understanding and informed participation as ways to transcend the individual human mind, our work has centered on developing the EDC as a research prototype. The EDC is based on the convergence of various systems (as shown in Figure 1) to create an integrated environment capable of addressing the following specific challenges: (1) How can we bring a variety of aspects (social, cultural, physical, virtual) together to support the creation of shared understanding [Resnick et al., 1991]? (2) How can we create co-evolutionary environments, in which stakeholders change because they learn, and in which systems change because stakeholders become co-developers and engage in end-user modification and programming [Mackay, 1992]? (3) How can we create intrinsically motivating computational environments and open systems, in which stakeholders feel in control and accept the role of active contributors rather than passive consumer [Fischer, 1998]? (4) How can stakeholders incrementally construct domain models that do not a priori exist but instead are socially constructed over time by communities of practice [Lave, 1988]?




Figure 1: The EDC as a convergence of systems

Figure 2 shows the current realization of the EDC environment. By using a horizontal electronic whiteboard (referred to in the scenario as the *action space*), participants work “around the table,” incrementally creating a shared model of the problem. They interact with computer simulations by manipulating the three-dimensional, physical objects that constitute a language for the domain [Ehn, 1988]. The position and movement of these physical objects are recognized by means of the touch-sensitive projection surface. In Figure 2, users are constructing a neighborhood through the use of a physical language appropriate for the problem by placing objects. This construction becomes the object through which the stakeholders can collaboratively evaluate and prescribe changes in their efforts to frame and resolve a problem. In the upper half of Figure 2, a second vertical electronic whiteboard (dubbed the *reflection space*) presents information related to the problem-at-hand for exploration and



Figure 2: The Current Prototype of the EDC

extension. In the figure, a user is filling out a survey constructed from the model presented in the action space. The results of this survey are stored (for future exploration) and are also fed to the simulation, where the ramifications of the decisions specified in the survey can be explored.

A Scenario: Creating Shared Understanding through Collaborative Design

The most mature EDC prototype application is one developed to support citizens in designing a transportation system for their neighborhood. Although this prototype has not yet been used in a real-world setting, its design has been shaped by the feedback we have received during participatory design [Ehn & Löwgren, 1997] and demonstration sessions with transportation domain experts, community activists, and peers within the HCI community. We describe in the following scenario, based on actual problem situations in the City of Boulder, how the EDC could be used across multiple design sessions to support citizens in planning a new bus route to service their community. In doing so, we will focus on three important facets of the EDC: (1) how participants interact with the system, (2) how they explore complex design problems, and (3) how they collaboratively construct new knowledge and incrementally create a shared understanding as they frame and resolve these problems.

A Neighborhood's Transportation Needs. Traffic and parking have become major problems in and around the city. A local neighborhood group, recognizing their area's contribution to the overall situation, has approached the city transportation planners to develop alternative transportation solutions. Current low-frequency bus routes have had little impact other than to generate comment that the large buses are frequently empty. To study the problem and to open a broader dialog with the neighbors, city planners convene an open meeting of various stakeholders (the concerned neighbors, transportation planners, and other city officials) using the EDC Urban-Planning application.

Creating a Language of Objects and Interacting with the System. The EDC Urban-Planning application uses a model and language that allows users to interact with various phenomena relevant to transportation planning. Objects in the system have both physical and computational representations. The physical objects represent language elements from the problem domain—in this setting, a language of colored blocks represents elements important to land use and transportation, such as residences, schools, shopping centers, parks, roads, buses, cars, and bus stops. These are linked to their computational representations through the EDC. The behavior and attributes of the language elements are represented in the computational objects, which can be defined or modified using an end-user, visual programming substrate called Visual AgentTalk [Repenning & Ambach, 1996].

This specific model, previously seeded by a collaboration of domain experts and citizens and evolved through actual use, simulates the dynamics of a bus route and contains specific information pertinent to the City of Boulder, such as population density, walking distances to bus stops, and waiting times. In this way the system is seeded with domain knowledge that will help guide the citizens as they explore transportation issues in their neighborhood. The seed provided will continue to grow and evolve at the hands of the citizens through its use [Fischer & Scharff, 1998b].




Figure 3: Placing Objects in the Action Space




Figure 4: Drawing Roads

The stakeholders begin framing the problem context—collaboratively constructing a description of their neighborhood by placing appropriate physical pieces on the interaction surface. The participants select objects from a palette, and the touch-sensitive tablet recognizes the location of the objects as they are placed. When neighbors place physical objects on the board the EDC creates a computer representation, which instantiates the object's behavior and default attributes (see Figure 3). The neighbors then create

roads to connect the different elements of their neighborhood. In this example, a road has behavior that automatically adds curves and intersections as necessary (see Figure 4).

Exploring Complex Problems. In addition to group construction, the EDC supports collaborative problem solving. Once the model is built, the neighbors indicate which way they travel to various destinations by using electronic markers to connect homes, schools, and shopping centers (see Figure 5).




Figure 5: Specifying Travel Preferences

This allows them to identify where transportation demands are heaviest and lightest and guides informed decisions regarding bus route placement. After the bus route is in place, the EDC's computational model simulates the behavior of the constructed bus system. In this way, the EDC supports the exploration of complex systems in a dynamic and contextualized manner. It is dynamic because the simulation shows how the model behaves (e.g., the neighbors see how the bus *travels* along the route), an attribute that our earlier physical decision-support games did not afford

[Arias et al., 1997]. It is contextualized because it is situated in a real task that the participants encounter and because the resolution grows out of the shared understanding that emerges as neighbors begin to better understand each other's perspectives regarding the neighborhood as they construct the model from *their* own understanding of *their* neighborhood. This is important because each participant may come to the table with different, often tacit [Polanyi, 1966] concepts about the neighborhood.

In some design sessions, neighbors model their neighborhood and design a new bus route. In such cases, the activity is constrained by existing infrastructure, such as roads and buildings.




Figure 6: Retrieving Constructions from the Reflection Space

In other design sessions, the bus route already exists and therefore is part of the existing problem context, so the task is not to define the location of the route, but rather to define the location of new bus stops along the route. The EDC stores existing constructions so that they can later be retrieved. For example, transportation planners discover that a particular bus route is underutilized so they set a meeting to get input from the community to see if bus stops can be better placed in order to increase the utilization of the route. To start the design

session, the participating citizens focus in on the neighborhood in question by selecting the proper section from an orthographic map (see Figure 6), which serves as the indexing mechanism for retrieval of existing constructions as well as a concrete context during the design session.

In this particular simulation the neighbors have modeled the use of the bus system for people traveling to school. The simulation presents different forms of information that may be important in understanding the transportation system to the participating neighbors. For example, the bus color represents whether it is empty (green), full (red), or in between. One of the neighbors is concerned that her workplace may be too far from the bus stop for her to use. By using the "walking distance" tool from the palette, she sees

that, in fact, her office is more than a 5-minute walk from the stop. The 5-minute walking radius is represented by the blue “X” marks (see Figure 7). She moves the bus closer to the center of the industrial park so that its 5-minute radius better covers the area.




Figure 7: The Color of the Bus and X's Surrounding the Bus Stop Enable Participants to Visualize Information in the Simulation

Continuing to run the simulation, another participant notices that the bus remains green-colored most of the time, which indicates that the bus is underutilized. After studying the model, he realizes that there is no bus stop serving the major residential area of the neighborhood. The neighbors discuss the problem and agree on a location for a new bus stop to service the residential area. When they add the bus stop to the model, the EDC creates a screen capture of the model and then automatically generates and displays a Web-based survey in the reflection space, which solicits ridership behavior data from the participants (see Figure 8).

The group discusses the survey and answers the questions in a way that best represents their behavior (e.g., they specify how long they will wait for the bus, based on various conditions). The ridership behavior data collected by the survey is then used to parameterize the simulation with the neighbors’ preferences. While the simulation runs, each bus stop keeps track of how often the bus arrives. If the bus does not arrive often enough, based on the survey information, then people waiting at the bus stop will drive their cars instead of using the bus.

The group discusses the survey and answers the questions in a way that best represents their

Learning on Demand. As cars begin to emerge in the simulation, the EDC displays information about this event in the reflection space. This signals a breakdown [Fischer, 1994c] in the model that they have constructed (i.e., people are not using the bus because it is not arriving frequently enough to suit their needs). The structure of the reflection space, implemented as an extensible Web site, provides an avenue for the neighbors to explore and to reflect upon the ramifications of the design choices that they have made in the action space (see Figure 9). This is a form of critiquing mechanism [Fischer et al., 1998], which links relevant information to the current breakdown. In the reflection space they see a brief description of the issue, which serves to ground their reflection to the emergent phenomena observed in the action space.

Next, the group explores and learns about the facts supporting different sides of the issue. One of the neighbors, whose priorities lean toward environmental responsibility, points out that increased car use can lead to increased air pollution. She supports her argument with the information she finds in the reflection space. Another participant stresses the convenience and flexibility of taking her car to work. If the buses arrived more frequently, she might consider taking the bus more often. The factual resource material found in this section of the reflection space provides a foundation from which the group members will be able to form their own opinions.

Once the participants have learned some of the objective facts surrounding the issue, they begin to explore the subjective opinions of other members of their community. Here they will begin to use the objective facts they have previously learned to evaluate others’ opinions. The participant whose priority was convenience points out a news article with the headline “SKIP Bus Causes Some Stress.” It describes the plight of one particular motorist who often finds herself caught behind a frequently stopping bus as she travels to and from work. Through this exploration, the opinions of some of the neighbors in the group begin to reformulate.




Figure 8: A Web-Based Survey Allows Participants to Parameterize the Bus Stop's Attributes

Learning the facts and understanding the thoughts and feelings of the community helps to provide a rich context for the current breakdown. This information delivered out of context would be of little use. These neighbors would not have considered exploring this information in the abstract, but in the context of the problem they are trying to resolve it has greater meaning and value. By delivering information that is situated to the activity taking place around the table, the participants can process the information more deeply than if it were encountered or delivered in a decontextualized manner.

Having a better understanding of the issue, the neighbors revisit the model they have constructed. The environmentalist of the group decides the solution is to add a few additional buses to the route. The group sees that this all but eliminates car use.

Meanwhile, the EDC continuously calculates the cost of the bus route, and one of the neighbors notices that they have just tripled the cost! Seeing this information, they all agree that this solution is not feasible.

Constructing New Knowledge. Faced with this dilemma, one of the neighbors recalls that some cities have implemented a light-rail train system to accommodate their citizens. He wonders if this would be a cheaper solution to their problem and asks if anyone knows anything about this alternative. None of the neighbors have any direct knowledge about light rail so they post a question to a discussion forum in the reflection space, which is juxtaposed with the domain knowledge that was supplied by the transportation experts. By doing this, the group documents an open issue that they would like to resolve before they meet again next week.

As the design session comes to a close, the group members agree to explore the light-rail question on their own before they meet again. While at home, each searches the Web for information on light-rail systems. As they find information that supports their individual perspectives, they add comments and URLs as responses to the original light-rail question posted by the group during the previous meeting (see the "Tell us what you think" pane in Figure 9). This allows members to collect information that will




Figure 9: The Reflection Space Combines Domain Expert Knowledge with an Extensible Discussion Forum to Provide a Mechanism for Learning on Demand and the Construction of New Knowledge

support their position at the next meeting. Through the face-to-face discussion that took place around the table and comments that each member posted to the discussion forum, the group members begin to understand each other's positions more clearly, and in some cases the perspectives of the members begin to converge.

Although the transportation domain experts seed the information space, the space is extensible and will evolve to embody the knowledge and opinions of the citizens using the system. This allows the group to develop a shared understanding of the problem and of each other.

The Conceptual Principles behind the EDC

The EDC effort is based on our collective prior work in the diverse fields of HCI and urban planning. From the HCI perspective we have engaged in the cultivation of conceptual frameworks and the creation of computational systems, such as domain-oriented design environments [Fischer, 1994a]. The urban-planning contributions include the notions of participation [Arias, 1989;

Arias, 1984], and the development of physical models and physical-simulation games [Arias, 1996] as decision support tools to empower citizens in the framing and resolution of complex planning problems, which by nature exist in a context of change and conflicting objectives [Arias, 1995; Jung et al., 1995].

Insights from these earlier efforts indicate that supporting a collaborative design process that includes both reflection and action requires a framework that can:

- deal with a set of possible worlds effectively (i.e., support exploration of design alternatives) to account for the fact that design is an argumentative process in which the goal is not to prove a point but instead to create an environment for a design dialog [Ehn, 1988; Simon, 1996];
- incorporate an emerging design in a set of external memory structures [Bruner, 1996], and record the design process and the design rationale [Fischer et al., 1996];
- generate low-cost, modifiable models that assist stakeholders in creating shared understanding by engaging in a “conversation with the materials” [Schön, 1983];
- use simulations to engage in “what-if” games and to replace anticipation of the consequences of our assumptions by analysis [Repenning & Sumner, 1995];
- make argumentation serve design [Fischer et al., 1996] and support reflection-in-action [Schön, 1983] by integrating action and reflection spaces; and
- introduce the notion of a common language of design by integrating physical objects with virtual objects [Arias, 1996].

The EDC Architecture

The architecture of the EDC, guided by insights and understanding gleaned from previous system design efforts [Fischer, 1994a], reflects the emerging requirements of the underlying evolutionary processes and application-domain support that it embodies. Initial efforts with the EDC focused solely on application to transportation planning in Boulder. Both generalizations and specializations of the architecture has occurred as we have moved to other settings and domains.

The architecture that has grown out of this process consists of three layers represented in Figure 10:

- The top layer realizes the conceptual principles discussed in the previous section through the integration of multiple system components as described in Figure 1.
- The middle layer represents different application domains (e.g., the domain of urban planning as illustrated in the scenario) and the domain objects specific to each.
- The bottom layer contextualizes this architecture with specific information for a particular application (in our case, transportation planning in our own community, the City of Boulder).




Figure 10: The EDC Layered Architecture

This diagram and process are based on the observation that socio-technical systems are never simply “instantiated” in any given individual situation. Each case has its own particularities, which are revealed and responded to as activity is produced for that case [Henderson, 1998].

The Integration of Action and Reflection

One of the primary theories behind the EDC is that (as discussed in the “Challenging Problems” section) people act until they experience a breakdown, this breakdown leads them to reflect upon their activities, and in this context they explore information spaces associated with the activity (Schön’s theory of reflection-in-action).

The EDC parallels this theory by providing support for action, support for reflection, and mechanisms that blend the two activities. In general, *action activities* take place on and around the horizontal table in Figure 2 and on the left side of Figure 1, through collaboration using a physical and computational model appropriate for the particular application domain. The scenario presents such a model for the EDC-Urban domain (providing a simulation with physical game pieces appropriate for modeling urban transportation problems), and uses context-dependent information (such as aerial photographs) for the specific application. *Reflection activities* are supported by the vertical white-board in Figure 2 and on the right side of Figure 1, through the capture, creation, presentation, and modification of hypermedia information [Moran & Melle, 1998]. This provides a portal to a dynamic, user-extensible, emergent Web-based information environment. In the scenario, the priority specification, maps, previous constructions, surveys, and critic information are stored and made available to support reflection activities.

The EDC supports ways to blend together these two aspects of reflection-in-action (see Figure 11). Critics




Figure 11: Blending Action and Reflection

[Fischer et al., 1998] are active agents that observe the collaborative construction and link to information relevant to the constructed artifact, such as when the cars begin to appear in the action space in the scenario. In the reflection space, there are generic mechanisms to capture and manipulate Web-based information to contextualize the design activity, as shown with the orthographic map and stored constructions in the scenario.

Both of these forms of activity, along with the mechanisms that support their integration, help make information relevant to the task at hand, support the interaction of multiple stakeholder perspectives, and draw on the various strengths that each brings to the task, resulting in collaborative exploration of the knowledge and construction of shared understanding about the problem.

It is important to understand that there is no strict dividing line between these two types of activity. Reflection can occur directly within the context of action, for example, when feedback from a simulation based on one action triggers several “what-if” actions by a participant. The participant then can explore and understand the consequences of decision options without resorting to a separate information space to explain the issue. Action can also take place within the information spaces that support reflection as new information is constructed, externalized, and reorganized. The most important contribution of the EDC is the synergy that is created between the action and reflection activities.

The EDC as an Open System

To support designers in framing and resolving their own problems, the EDC needs to support a dynamic evolving problem context. Exemplifying open principles is important in addressing open-ended problems and collaborative creation of shared understanding in the EDC. In a domain such as transportation planning, no system can completely subsume all information needed to solve a problem. An essential goal of the system is to provide a shared representation that all participants can extend when the need arises. In fact, the extension process itself may play an important role in creating shared understanding by supporting the collaborative activity of extending the realization of the problem.

On the technical level, all of the components used to create the EDC environment are designed to be extensible by users. In the action space, the physical language provides an initial tool to describe a problem, but users might choose to add new objects to the language to represent new kinds of objects. In our current models, which use colored blocks, one might introduce a new object with a different color or shape. The corresponding computational model can be modified as situations arise. We are currently using AgentSheets [Repenning & Sumner, 1995] and the associated end-user programming language Visual AgenTalk [Repenning & Ambach, 1996] as substrates for building simulations within the EDC. AgentSheets and Visual AgenTalk allow users to quickly add or change the objects that make up a model and experiment with changes they make to the computational model. The dynamic information spaces used in the reflection space are designed to allow users to extend information. We currently use the DynaSites substrate [Ostwald, 1999], a tool for making evolvable Web-based information spaces.

The scenario demonstrated many levels at which extensions might take place in the EDC. Users can add information to a problem situation by entering new or linking existing content into the evolving Web documents maintained by the reflection space. Filling out a survey extends the information available to the computational simulation, in this case altering the parameters of a specific bus stop. Additional parameters are available for every element in the simulation. If existing parameters do not capture the kind of modification that users wish to make, users can change the behavior of individual objects in the simulation.

Although providing support for modification at all levels is an important step toward making the EDC an open system, merely providing opportunities for extension is not enough to truly support open

evolution. One of the major challenges of the EDC is to provide both a technical and a social context appropriate for evolution [Fischer & Scharff, 1998a; Raymond, 1998]. For example, in order to add pieces to the simulation, users currently need an understanding of the whole model, which requires a good understanding of modeling with Visual AgenTalk to modify the simulation. Creating a model for extension that is tailored to a given situation and created with an understanding of the background of the users is an important future direction. One of the major future challenges for evolution in the EDC is not simply to make the system evolvable at all levels, but to provide a use context in which evolution can be captured through collaborative activity using means that are appropriate for the problem and target audience.

Assessment

In the HCI community, the late 1980s and early 1990s marked the years of the novice, when conventional assessment (e.g., experimental psychology techniques using laboratory tasks) worked well [Bannon, 1995; Thomas & Kellogg, 1989]. Since the mid-90s, HCI has increasingly moved more into the era of skilled domain workers. The activities and processes that we want to support with the EDC, as argued in the introduction, take months, years, and decades. We must account for the rich context in which design takes place and create situations grounded in practice. As a result, our goal for the assessment of the EDC effort is to transcend the laboratory and analyze and evaluate our environments in real-world settings. While the EDC as a whole has not yet been put into broad practice and evaluated, we have had considerable experience with the assessment of essential parts of our system.

Assessment in Design and Practice. In our approach to design, assessment is viewed not as the endpoint of a waterfall model but as a process integrated into design and practice. The design of the EDC is based on assessment of our own prior work (as discussed in [Arias et al., 1997]) as well as a study of the strengths and limitations of other theoretical work, approaches, and systems, including ubiquitous computing [Abowd et al., 1998; Weiser, 1991; Weiser, 1993], collaboratories [Erickson et al., 1999; Olson & Olson, 1997], and “Roomware” [Streitz et al., 1994; Streitz et al., 1999]. This is an ongoing activity throughout the design process, not just the starting point for our investigations.

Crucial insights from our prior work that have laid the groundwork for our design of the EDC are based on our use of physical simulations applied to actual community design with specific neighborhoods (e.g., [Arias, 1996]). These insights, along with other efforts on how we can create representations that can be shared and understood by all stakeholders, have indicated that physical objects are critically important. This has been borne out at two levels. First, the direct, naïve manipulability of physical objects is important for special groups who may not be well versed in technology. Second, we have seen the importance of the innate understanding that comes from manipulation of physical objects.

Assessment through Participatory Design. By involving communities of practice in the design of EDC domain prototypes, we have gained considerable insight into how things are (settings, cultures), how they are done (processes, organizations), why they are the way they are, and how they are limited by current practice.

Our work in this area has focused on participatory design efforts based on numerous joint design sessions with the Boulder County Healthy Communities Initiative and the Regional Transportation District in the Denver-Boulder County Region of Colorado. We have gained critical insights into the design and development of the EDC through these interactions. These include:

- the importance of being able to represent multiple perspectives [Stahl, 1999] of a problem;
- the need to support learning as a shared, collaborative activity—particularly in the context of bridging these multiple perspectives;
- the potential use of the EDC to provide support for democratic and social processes; and
- the need to support interaction and reflection both “around the table” as well as “beyond the table.”

Assessment of Open Systems and Emerging Applications. The emergence that takes place in an open system will not take place within the first few days or weeks of use—this makes an experimental psychology approach of hiring subjects and measuring their interaction with the system impossible. We need to understand the long-term use of a system by owners of problems engaged in the cultivation of a rich repertoire of personally and socially meaningful artifacts. We do not expect all users to become

Visual AgenTalk programmers or to be interested in making radical changes to the system. Users' contributions will depend on the perceived benefit, which involves the effort needed to make changes and the utility received for effecting changes.

Assessment of the Effectiveness of Interaction Techniques. Although low-level human-technology interaction techniques [Newell & Card, 1985] are not the primary focus of our work, nonetheless, they are an important aspect of designing for the activities we want to support. The current touch-screen realization of the action space implicitly creates a turn-taking and modal interaction. We have observed breakdowns when two users try to place objects at once (causing the system to draw objects between the two placements) or place objects that differ from the currently selected object (that is, a user tries to place a home, but because the system is in "school mode," a school gets placed in the simulation instead). People unfamiliar with the technology get confused at these violations of the assumptions they have made about the technology. As we continue to develop the EDC, we will evaluate the effectiveness of interaction through analysis of the breakdowns and successes of the technology through design, demonstration, and use activities.

Future Work

Assessment of Support for the Creation of Shared Understanding. Supporting "around-the-table" interaction and contextualizing information in design activities are critical elements in creating shared understanding. It is important to discover which social situations are more conducive to the creation of this shared understanding. For example, important aspects to study include determining the utility of a trained facilitator, the efficacy of participant facilitators, and the effect that such interventions would have on "putting the owners in charge" [Fischer, 1994b]. By analyzing how the EDC is utilized during design activities, we will assess the social and technical dimensions of how shared understanding can be created. An important issue for assessment will be to track long-term effects of the design processes upon the design community as well as to evaluate the products of design.

This assessment will take place against a backdrop of experiences with organizational memories and collaborative work that have exposed two barriers to capturing information: (1) individuals must perceive a direct benefit in contributing to organizational memory that is large enough to outweigh the effort [Grudin, 1989]; and (2) the effort required to contribute to organizational memory must be minimal so it will not interfere with getting the real work done [Carroll & Rosson, 1987].

Use of the EDC in Actual Work Situations. Although we have gained a great deal of insight into the design and effectiveness of our approaches through the integrated activities we have already employed, there are still critical perspectives to be gleaned from deployment and study of our systems in use contexts. We will utilize insights from activities such as ethnographic methods "in the wild" [Hutchins, 1994], studies of everyday activities [Nardi & Zamer, 1993], and analysis of conversational interaction [Goodwin & Heritage, 1990; Jordan & Henderson, 1995].

Beyond Binary Choices. By arguing for the desirability of supporting people as designers, we want to state explicitly that there is nothing wrong with being a consumer and that we can learn and enjoy many things in a consumer role. It is a mistake to assume that being a consumer or being a designer would be a binary choice—it is rather a continuum ranging from passive consumer, to active consumer, to end-user, to user, to power users, to domain designer, to medium designer. Problems occur, for example, when someone wants to be a designer but is forced to be a consumer or when being a consumer becomes a universal habit and mindset dominating one's life completely. We claim that the HCI community should not be content with either (1) restricting its efforts to the user interface or the computational aspects of HCI, or (2) reflecting and evaluating designs developed by other communities (e.g., the groups who give us 500 TV channels or artifacts over which we have no control). The HCI research community should not confine itself to a consumer role in the process of shaping our future knowledge society [Drucker, 1994] in which they focus solely on some technical issues in the context of a world defined by others.

Conclusions

The EDC is a contribution toward creating a new generation of collaborative human-computer systems that address and overcome current limitations of human-computer interaction. It shifts the emphasis

away from the computer screen as the focal point and creates an integrated environment in which stakeholders can incrementally create a shared understanding through collaborative design. It is an environment that is not restricted to the delivery of predigested information to individuals; rather, it provides opportunities and resources for design activities embedded in social debates and discussions in which *all* stakeholders can actively contribute rather than being confined to passive consumer roles.

HCI research and development have made very important contributions over the last decade. The HCI community has acquired a broad understanding of creating computational artifacts fitting better human capabilities and needs by creating theories and innovative systems [Helander et al., 1997]. To take the next step forward, the HCI community should accept the challenge of rethinking computational media in broader contexts. Our claim is that computational media can have an impact on our individual lives and our societies similar the fundamental change from oral to literal societies brought about by the introduction of reading and writing. The true contribution of computational media may be to allow all of us to take on or incrementally grow into a designer role in areas that we consider personally meaningful and important such that we are motivated to expend the additional effort. The future of HCI lies in realizing that what we can build is more limited by our imagination, our ability to discover, and our ability to envision than by our system development limitations.

Acknowledgments

The authors thank Taro Adachi, Josh Emhoff, Rogerio dePaula, Christine Giger, Volker Jung, Shigeru Kurihara, Kelli Murphy, Jonathan Ostwald, Alexander Repenning, Kurt Schneider, Stefanie Thies, and Jessica Witter for their help in conceptualizing, implementing, and supporting the EDC as well as with the preparation of this paper. This work is supported in part by NSF grants CDA-9529549, REC-9631396, and IIS-9711951 and by funding from PFU, Inc, Tokyo.

References

- Abowd, G. D., Atkeson, C. G., Brotherton, J. A., Enqvist, T., Gully, P. A., & Lemon, J. (1998) "Investigating the Capture, Integration, and Access Problem of Ubiquitous Computing in an Educational Setting," *Proceedings of CHI '98*, pp. 440-447.
- Arias, E. (1989) "The Contingent Nature of Participation and Housing Research.," *Guru Nanak Journal of Sociology—Special Housing Policy Issue*, 10(1-2), pp. 81-99.
- Arias, E. G. (1984) "Resident Participation in Public Housing: A Conceptual Approach," In Vogt & Mickle (eds.), *Part 1: Geography-Regional Science*, ISA, Research Triangle Park, NC, pp. 131-138.
- Arias, E. G. (1995) "Designing in a Design Community: Insights and Challenges," In G. M. Olson & S. Schuon (eds.), *Proceedings of DIS'95 Symposium on Designing Interactive Systems: Processes, Practices, Methods, & Techniques*, ACM Press, New York, pp. 259-263.
- Arias, E. G. (1996) "Bottom-up Neighborhood Revitalization: Participatory Decision Support Approaches and Tools," *Urban Studies Journal—Special Issue on Housing Markets, Neighborhood Dynamics and Societal Goals*, 33(10), pp. 1831-1848.
- Arias, E. G., Eden, H., & Fisher, G. (1997) "Enhancing Communication, Facilitating Shared Understanding, and Creating Better Artifacts by Integrating Physical and Computational Media for Design," In *Proceedings of DIS'97: Designing Interactive Systems: Processes, Practices, Methods, & Techniques*, ACM, New York, pp. 1-12.
- Arias, E. G. & Schneider, K. (1999) "Decision Support for Wicked Planning Problems," *Journal of Simulations and Games (under revision)*.
- Bannon, L. J. (1995) "The Politics of Design: Representing Work," *Communications of the ACM*, 38(9), pp. 66-68.
- Brown, J. S., Duguid, P., & Haviland, S. (1994) "Toward Informed Participation: Six Scenarios in Search of Democracy in the Information Age," *The Aspen Institute Quarterly*, 6(4), pp. 49-73.

- Bruner, J. (1996) *The Culture of Education*, Harvard University Press, Cambridge, MA.
- Carroll, J. M. & Rosson, M. B. (1987) "Paradox of the Active User," In J. M. Carroll (ed.), *Interfacing Thought: Cognitive Aspects of Human-Computer Interaction*, The MIT Press, Cambridge, MA, pp. 80-111.
- Drucker, P. F. (1994, November) "The Age of Social Transformation," *The Atlantic Monthly*, pp. 53-80.
- Ehn, P. (1988) *Work-Oriented Design of Computer Artifacts*, Almquist & Wiksell International, Stockholm, Sweden.
- Ehn, P. & Löwgren, J. (1997) "Design for Quality-in-Use: Human-Computer Interaction Meets Information System Development," In M. G. Helander, T. K. Landauer, & P. V. Prabhu (eds.), *Handbook of Human-Computer Interaction*, Elsevier Science B. V., Amsterdam, pp. 299-313.
- Eisenberg, M. & Makay, W. (1996) "Real Meets Virtual: Blending Real-World Artifacts with Computational Media," *CHI '96, Human Factors in Computing Systems*, pp. 159-160.
- Engelbart, D. C. (1995, August) "Toward Augmenting the Human Intellect and Boosting our Collective IQ," *Communications of the ACM*, pp. 30-33.
- Erickson, T., Smith, D. N., Kellogg, W. A., Laff, M., & Brander, E. (1999) "A Sociotechnical Approach to Design: Social Proxies, Persistent Conversations, and the Design of Babble," (to appear in) *Proceedings of CHI '99*.
- Fischer, G. (1991) "Supporting Learning on Demand with Design Environments," In L. Birnbaum (ed.), *International Conference on the Learning Sciences (Evanston, IL)*, Association for the Advancement of Computing in Education, pp. 165-172.
- Fischer, G. (1993) "Beyond Human Computer Interaction: Designing Useful and Usable Computational Environments," In *People and Computers VIII: Proceedings of the HCI '93 Conference (Loughborough, England)*, Cambridge University Press, Cambridge, UK, pp. 17-31.
- Fischer, G. (1994a) "Domain-Oriented Design Environments," *Automated Software Engineering*, 1(2), pp. 177-203.
- Fischer, G. (1994b) "Putting the Owners of Problems in Charge with Domain-Oriented Design Environments," In D. Gilmore, R. Winder, & F. Detienne (eds.), *User-Centered Requirements for Software Engineering Environments*, Springer Verlag, Heidelberg, pp. 297-306.
- Fischer, G. (1994c) "Turning Breakdowns into Opportunities for Creativity," *Knowledge-Based Systems, Special Issue on Creativity and Cognition*, 7(4), pp. 221-232.
- Fischer, G. (1998) "Beyond 'Couch Potatoes': From Consumers to Designers," In IEEE (ed.), *1998 Asia-Pacific Computer and Human Interaction, APCHI'98*, IEEE Computer Society, pp. 2-9.
- Fischer, G., Grudin, J., Lemke, A. C., McCall, R., Ostwald, J., Reeves, B. N., & Shipman, F. (1992) "Supporting Indirect, Collaborative Design with Integrated Knowledge-Based Design Environments," *Human Computer Interaction, Special Issue on Computer Supported Cooperative Work*, 7(3), pp. 281-314.
- Fischer, G., Lemke, A. C., McCall, R., & Morch, A. (1996) "Making Argumentation Serve Design," In T. Moran & J. Carrol (eds.), *Design Rationale: Concepts, Techniques, and Use*, Lawrence Erlbaum and Associates, Mahwah, NJ, pp. 267-293.
- Fischer, G., Nakakoji, K., Ostwald, J., Stahl, G., & Sumner, T. (1998) "Embedding Critics in Design Environments," In M. T. Maybury & W. Wahlster (eds.), *Readings in Intelligent User Interfaces*, Morgan Kaufmann Press, pp. 537-559.
- Fischer, G. & Scharff, E. (1998a) "Learning Technologies in Support of Self-Directed Learning," *Journal of Interactive Media in Education*, 1998(4), pp. www-jime.open.ac.uk/98/4.
- Fischer, G. & Scharff, E. (1998b) "Learning Technologies in Support of Self-Directed Learning," *Journal of Interactive Media in Education*, 1998(4).

- Galegher, P., Kraut, R., & Egidio, C. (eds.) (1990) *Intellectual Teamwork*, Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Girgensohn, A. (1992) *End-User Modifiability in Knowledge-Based Design Environments*, Ph.D. Dissertation, University of Colorado at Boulder.
- Goodwin, C. & Heritage, J. (1990) "Conversational Analysis," *Annual Review of Anthropology*, 19, pp. 283-307.
- Greenbaum, J. & Kyng, M. (eds.) (1991) *Design at Work: Cooperative Design of Computer Systems*, Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Grudin, J. (1989) "Why Groupware Applications Fail: Problems in Design and Evaluation," *Office: Technology and People*, 4(3), pp. 245-264.
- Grudin, J. (1994) "Groupware and Social Dynamics: Eight Challenges for Developers," *Communications of the ACM*, 37(1), pp. 92-105.
- Helander, M. G., Landauer, T. K., & Prabhu, P. V. (eds.) (1997) *Handbook of Human-Computer Interaction* (second, completely revised ed.), Elsevier Science Ltd., Amsterdam.
- Henderson, A. (1998, November & December) "An Interview with Austin Henderson by Kate Ehrlich," *Interactions*, pp. 36-47.
- Hutchins, E. (1994) *Cognition in the Wild*, The MIT Press, Cambridge, MA.
- Illich, I. (1971) *Deschooling Society*, Harper and Row, New York.
- Ishii, H. & Kobayashi, M. (1992) "ClearBoard: A Seamless Medium for Shared Drawing and Conversation with Eye Contact," In *Proceedings of CHI'92 Conference*, pp. 525-532.
- Ishii, H. & Ullmer, B. (1997) "Tangible Bits: Towards Seamless Interfaces between People, Bits and Atoms," In *Proceedings of ACM CHI 97 Conference on Human Factors in Computing Systems*, pp. 234-241.
- Jordan, B. & Henderson, A. (1995) "Interaction Analysis: Foundations and Practice," *Journal of the Learning Sciences*, 4(1), pp. 39-103.
- Jung, V., Arias, E. G., & Catherine, G. (1995) *Design of Integrated Decision Support Tools for Sustainable Environmental Planning*, (GRIS No. 94-3), Technical University of Darmstadt, Darmstadt, Germany.
- Lave, J. (1988) *Cognition in Practice*, Cambridge University Press, Cambridge, UK.
- Mackay, W. E. (1992) "Co-adaptive Systems: Users as Innovators," In *CHI'92 Basic Research Symposium*.
- Moran, T. & Melle, W. v. (1998) "Tailorable Domain Objects as Meeting Tools for an Electronic Whiteboard," *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work (CSCW'98)*, pp. 295-304.
- Moran, T. P. & Carroll, J. M. (eds.) (1996) *Design Rationale: Concepts, Techniques, and Use*, Lawrence Erlbaum Associates, Inc., Hillsdale, NJ.
- Myers, B. A. (1998, March). "A Brief History of Human Computer Interaction Technology." *ACM Interactions*, pp. 44-54.
- Nakakoji, K., Yamamoto, Y., Suzuki, T., Takada, S., & Gross, M. (1998) "From Critiquing to Representational Talkback: Computer Support for Revealing Features in Design," *Knowledge-Based Systems Journal*, 11(7-8), pp. 457-468.
- Nardi, B. A. (1993) *A Small Matter of Programming*, The MIT Press, Cambridge, MA.
- Nardi, B. A. & Zamer, C. (1993) "Beyond Models and Metaphors: Visual Formalisms in User Interface Design," *Journal of Visual Languages and Computing*(4), pp. 5-33.

- Newell, A. & Card, S. K. (1985) "The Prospects for Psychological Science in Human-Computer Interaction," *Human-Computer Interaction*, 1(3), pp. 209-242.
- Norman, D. A. (1990) "Why Interfaces Don't Work," In B. Laurel (ed.), *The Art of Human-Computer Interface Design*, Addison-Wesley Publishing Company, Reading, MA, pp. 209-219.
- Norman, D. A. (1993) *Things That Make Us Smart*, Addison-Wesley Publishing Company, Reading, MA.
- Olson, G. M. & Olson, J. S. (1997) "Research on Computer Supported Cooperative Work," In M. G. Helander, T. K. Landauer, & P. V. Prabhu (eds.), *Handbook of Human-Computer Interaction*, Elsevier Science B.V., Amsterdam, pp. 1433-1456.
- Ostwald, J. (1999) *DynaSites*, At <http://www.cs.colorado.edu/~ostwald>.
- PCSD—President's Council on Sustainable Development (1996) *Sustainable America: A New Consensus for Prosperity, Opportunity, and a Healthy Environment for the Future*, President's Council on Sustainable Development, Washington, D.C.
- PITAC—President's Information Technology Advisory Committee (1999) *Information Technology Research: Investing in Our Future*, National Coordination Office for Computing, Information, and Communications, Arlington, VA, At http://www.ccic.gov/ac/report/pitac_report.pdf.
- Polanyi, M. (1966) *The Tacit Dimension*, Doubleday, Garden City, NY.
- Raymond, E. S. (1998) *The Cathedral and the Bazaar*, At <http://earthspace.net/~esr/writings/cathedral-bazaar/cathedral-bazaar.html>.
- Repenning, A. & Ambach, J. (1996) "Tactile Programming: A Unified Manipulation Paradigm Supporting Program Comprehension, Composition and Sharing," In *Proceedings 1996 IEEE Symposium on Visual Languages*, IEEE Computer Society Press, Los Alamitos, CA,, pp. 102-109.
- Repenning, A. & Sumner, T. (1995) "Agentsheets: A Medium for Creating Domain-Oriented Visual Programming Languages," *IEEE Computer, Special Issue on Visual Programming*, 28(3), pp. 17-25.
- Resnick, L. B., Levine, J. M., & Teasley, S. D. (eds.) (1991) *Perspectives on Socially Shared Cognition*, American Psychological Association, Washington, D.C.
- Rittel, H. (1984) "Second-Generation Design Methods," In N. Cross (ed.), *Developments in Design Methodology*, John Wiley & Sons, New York, pp. 317-327.
- Rittel, H. & Webber, M. M. (1984) "Planning Problems Are Wicked Problems," In N. Cross (ed.), *Developments in Design Methodology*, John Wiley & Sons, New York, pp. 135-144.
- Roggoff, B., Matusov, E., & White, C. (1998) "Models of Teaching and Learning: Participation in a Community of Learners," In D. R. Olson & N. Torrance (eds.), *The Handbook of Education and Human Development*, Blackwell Publishers, Malden, MA, pp. 388-414.
- Scardamalia, M. & Bereiter, C. (1994) "Computer Support for Knowledge-Building Communities," *Journal of the Learning Sciences*, 3(3), pp. 265-283.
- Schön, D. (1992) "Designing as reflective conversation with the materials of a design situation," *Knowledge-Based Systems Journal, Special Issue on AI in Design*, 5(1), pp. 3-14.
- Schön, D. A. (1983) *The Reflective Practitioner: How Professionals Think in Action*, Basic Books, New York.
- Simon, H. A. (1996) *The Sciences of the Artificial* (third ed.), The MIT Press, Cambridge, MA.
- Snow, C. P. (1993) *The Two Cultures*, Cambridge University Press, Cambridge, UK.
- Stahl, G. (1999) "WebGuide: Guiding Collaborative Learning on the Web with Perspectives," In *Annual conference of the American Educational Research Association (AERA '99)*, Montreal, Canada.
- Streitz, N., Geißler, J., Haake, J., & Hol, J. (1994) "DOLPHIN: Integrated Meeting Support across LiveBoards, Local and Desktop Environments," *Proceedings of CSCW'94*, pp. 345-358.

Streitz, N. A., Geißler, J., Holmer, T., Konomi, S. i., Müller-Tomfelde, C., Reischl, W., Rexroth, P., Seitz, P., & Steinmetz, R. (1999) "I-LAND: An Interactive Landscape for Creativity and Innovation," In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI '99)*, ACM Press, New York, pp. 120-127.

Thomas, J. C. & Kellogg, W. A. (1989) "Minimizing Ecological Gaps in Interface Design," *IEEE Software*, 6, pp. 78-86.

Weiser, M. (1991) "The Computer for the 21st Century," *Scientific American*, 265(3), pp. 66-75.

Weiser, M. (1993) "Some Computer Science Issues in Ubiquitous Computing," *Communications of the ACM*, 37(7), pp. 75-83.