Touchable Interactive Walls: Opportunities and Challenges

Kelly L. Dempski[1], Brandon L. Harvey[1]

[1] Accenture Technology Labs, 161 N. Clark St.,
Chicago, IL 60601, USA
{Kelly.L.Dempski, Brandon.L.Harvey}@accenture.com
http://www.accenture.com/Techlabs}

Abstract. Very large, high resolution, interactive screens—also known as interactive walls—can be used to deliver entertainment and advertising content that is qualitatively different from what is available in television, kiosk, or desktop formats. At a sufficient resolution and size, the touchable wall can offer the engaging interactivity of full-fledged entertainment software, but on a scale that enables new kinds of public experiences. This paper describes some of the opportunities enabled by what we believe to be a new computing medium in its own right. We also describe some of the new design challenges inherent in this medium, together with suggestions based on our own approach to those challenges.

1 Introduction

Since the advent of movies, audiences have experienced entertainment presented on large screens, with actors and events often appearing larger than life. IMAX movies take the experience one step further; the content is specifically designed to encompass the viewer's whole visual field, enveloping them in a nearly immersive experience. Despite this feeling of immersion, of course, the viewer does not participate; the audience is passive.

On the other hand, video games and computer applications are highly interactive, but are usually relegated to the smaller physical formats of monitors or televisions. A few privileged users do have access to larger-format interactive screens such as SmartBoards, but these are rarely used as a medium for public entertainment. Where large interactive screens are deployed in public, they tend to feature the same sorts of content that would be found on a desktop computer, simply magnified onto a larger area. In one type of installation, the touch interface may be simplified to its barest essentials—pressing buttons—putting a limit richness of possible interactions; in the other direction, the user can control a pointer, making the interface is more flexible. But typically, few accommodations are made for the difference in circumstances between desktop computing and interactive touch screens. The same interface elements developed for a screen one foot wide—the menus and lists, the positioning of elements on screen and their relative proportions—are imported wholesale into a context where there is a much larger canvas, a different sort of pointer (the hand), and a different set of use cases.

We believe that large interactive displays will be effective only if the content is specifically designed to take advantage of the particular parameters of this medium. The overall scope of our research covers a range of applications, from large command and control screens, to group decision applications, to entertainment and advertising applications. The purpose of this paper is to cover the latter, and show the opportunities and challenges involved with creating content that is large in size, highly detailed, directly interactive for multiple users, and usable by the general public.

1.1 Defining the Interactive Wall

Our definition of an interactive wall includes both physical size and resolution. Just as a true IMAX experience requires both a large screen and a higher resolution film stock, high quality large format digital experiences demand both a large size and very high resolution—a lot of pixels in a small physical area. To create a large, high-resolution screen, people typically tile together a number of high-resolution display elements. Whether this is accomplished using front or rear projection, or by stacking up special monitors with very thin bezels, two problems immediately arise. The first is the number of video inputs—what will drive all of the video ports? The second is the sheer number of pixels involved—what will compute all the screen contents? There are few if any graphics cards that can generate a canvas many thousands of pixels across, with real-time performance.

There are several approaches to these challenges, ranging from specialized signal fusion hardware [1] to low level software schemes for distributed rendering [2]. Because of limitations of these approaches, we used neither and created a higher level distributed rendering scheme (described in [4]) that allows us to distribute a very large high resolution application across a network of heterogeneous devices. Although we can scale much higher, our current screen is comprised of 8 1024x768 tiles running on 4 PCs connected with our software. This effectively creates a virtual PC capable of running at 4096x1536 resolution on a 10ftx4ft screen with very high 2D and 3D graphics performance and a low cost. This resolution will be mentioned later when we compare desktop applications to potential wall applications, although our system supports different configurations and different display resolutions.

Our approach is focused on supporting interactive content. To accomplish this, we built a very high resolution touch system, based on cameras, that senses the entire screen with a high sampling rate (greater than 120 Hz) and a precision finer than a fingertip. This setup allows us to control the virtual application with the same responsiveness one would expect from a mouse. More recently, we have developed the system's ability to detect and resolve multiple simultaneous touch inputs. We find this very exciting because it allows us to begin creating interfaces that can be used either separately or cooperatively. This aspect can be very important in public spaces where multiple people may want to interact with content simultaneously.

To date, related projects have focused on the technological infrastructure—without adapting content to fit the new form factor—or they have addressed size and resolution issues independent of interactions [3][5]. We have found that the combination of all of the features of an interactive wall, working in concert, demands new approaches to content creation and application design.

The combination of the aforementioned features comprises what we believe to be a new medium with new capabilities. In the remainder of the paper, we'll discuss the application opportunities created by the wall medium; then we will describe the related challenges and some approaches to meeting these challenges.

2 Opportunities for Interactive Walls

In this section, we will discuss how large format interactive walls are qualitatively different from standard kiosks (which facilitate user interaction) and large passive displays (which emphasize a visual experience). We will discuss each of these elements separately to better explain the value of each. However, it is important to recognize that the best applications will be the ones that blend all the elements to create rich interactive experiences.

2.1 Touch: Direct Experiences

When you interact with a computing system you become engaged in an entirely different way than when you passively absorb information from a screen or a sign. This is one of the primary reasons why computer entertainment exists, and it is also a principle that has been put to work many times in, for example, educational software.

An interactive wall that can be touched by anyone, using no special tools or skills, gives users the chance to shape an experience as they might with a computer, but on a scale similar to that of physical experiences. Interactive walls can reward sustained attention, like computers do, but they can also support a more casual interaction pattern than desktop computers do, because they forgo the small, delicate controls and the inert sitting posture. On an interactive wall, even very broad physical input—such as waving a hand over something—can create large, dramatic results, results that are instantly shareable with others because of the public nature of the display.

An interactive wall, if it allows multiple simultaneous touch inputs, may focus on bringing people together into entertaining activities in which they compete or cooperate (a life-size game of Pong); or it might allow users to interact separately (a music store could have a wall that functions as a row of "listening stations"). Interacting with users who are distant, perhaps standing at other walls, could enlarge the social dimension of the wall.

Advertisers may use interactive walls as a way to get people more involved with a campaign or a product. The content might involve: product information; games or prizes; demonstrations of a product (users could interact with virtual doubles of real products such as music players or cars); offering special functionality or "Easter eggs" to existing customers (of, say, a particular wireless device); and so forth.

2.2 Large Size: Life Size Experiences

A screen so big that it takes up a user's entire field of view provides one of the most immersive experiences we can create, short of virtual reality goggles. IMAX screens produce this effect for users who are many yards away, by deploying a truly enormous screen. In the case of an interactive wall, which is merely large, the enveloping nature of the experience is a byproduct of how close the primary users are.

Some applications benefit from a large size presentation simply because it makes more sense to present certain objects at a lifelike size. For example, human beings have seldom been represented before in an interactive medium at a true scale. With a wall, the user can be face to face with someone who appears to be standing right in front of them. (In fact, with the right setup, it is possible to stage correct eye contact between two remote users.) Because of factors like these, applications such as teleconferencing, distance learning, or remote business presentations may be more successful at wall-size than they have been in other form factors.

Conveying aesthetic details or portraying relationships in space have presented challenges in the consumer-oriented visualization area. But since even large objects such as cars can appear at near life size on a sizable interactive wall, there may be new viability for some forms of "virtual product showroom." For example, a home furnishings retailer might find a competitive advantage in offering customers an interactive experience of how their home would look, decorated and furnished differently, in a life-size virtual mockup.

The number of and kind of users anticipated to approach the wall will guide the design of the wall's size and shape. If we expect dense crowds around the wall, blocking lines of sight, we may want to make the display some feet taller than user's reach. This leaves space to display content, once the crowd has packed around the display. Conversely, a wall intended for children may be built much shorter, and set closer to the ground.

2.3 High Res: Life-Like Experiences

Data density is important not just in scientific applications, with their complex information sets, but for entertainment and advertising as well: text elements, detailed photographs, generated graphics, images of fabrics and textures, and human faces all benefit from detailed presentation. A wall might feature incredibly detailed scenes, thousands of pixels across, whether recorded, artificially generated, or streamed live from another location. Depending on the quality of the content, things onscreen may appear with near-ocular visual fidelity, as a result of a resolution that is multiples higher than what is available through DVD or HDTV.

Designers may take high resolution into account to create "stations" for individual users, within the larger canvas of the wall. When multiple users are standing shoulder to shoulder, interacting simultaneously, high pixel densities let us fit rich visual information into a manageably small space in front of each user.

At the same time, the effects of user interactions should be designed to be large enough so that secondary users, standing behind the direct users, can share the experience—suggesting actions, gleaning information, or simply being entertained or informed.

It is important to note that unlike in the case of most other screens, the person touching an interactive wall may be unable see the entire surface, except at a very steep viewing angle, and not all at the same time. Those standing farther back, of course, have a different experience: they see the whole wall at once, and for them, the users and their actions on the wall actually function as a part of the display—part of the experience. This is a feature that can inform design choices; an application (such as a puzzle) might involve a number of localized tasks for primary users to carry out, while users standing farther back have a special role because they can see the "big picture".

2.4 Multi-User: Shared Experiences

The user's interactions, and the kinds of applications a wall can offer, take on a different light when we consider that, unlike most other public computer systems, which at least try to create a certain privacy, wall interactions are completely, unapologetically public. Given the size, brightness, and the geometry of the wall, and given the fact that the users interact with relatively large motions of their hands and arms (rather than small motions of the fingers and wrists), both the input and the output of the wall are manifestly obvious to everyone around. Wall applications will automatically be spectacles.

Laptop computers, ATMs, and airport kiosks are used in public places, but in a semi-private fashion (at least hopefully), because of the personal information they convey. Wall applications, by contrast, will typically avoid any deployment of personal data. In fact, this will be one the main attractions of the medium: as with pick-up basketball games, strolling on a boardwalk, and karaoke, watching and being watched are both acceptable forms of participation in the public experience. Walls will be appropriate wherever crowds are appropriate, or wherever one wishes to try to draw a crowd.

We can define three criteria for a successful wall application: (a) it must attract—and then entertain, sell to, or educate—interacting users (b) it must keep those users happy about performing their interactions in front of an actual or potential public (c) it must also perform an entertainment, selling, or education function for onlookers.

By the lights of these criteria, interactive walls applications share more qualities with blackboards than they do with personal computers. A blackboard in a hallway is a place to write, but also to be seen writing—people occasionally write there for their own understanding, but usually on behalf of (actual or potential) others as well. It is the sense of being seen that consciously or unconsciously governs what is written. A similar sensibility can be expected to influence users' behavior as they interact with various wall applications.

The arcade game provides a successful example of public computing, and one that satisfies the same three criteria: the user is entertained by the interaction; she is conscious of, and comfortable with, her role in creating a potential or actual public spectacle; and others are free to look on. These three factors considered together define the value of the device.

2.5 Connectedness: Dynamic Experiences

Some of the examples above describe instances where the wall supersedes a non-interactive medium, and the dynamic nature of the display is used to create engaging interactions and spectacles. However, the content is also dynamic in another sense of the word; the contents of a networked display can be changed or updated more easily than standard signage. This is very important to advertisers who might want to change messages depending on audience, time of day, or business conditions.

The same might be accomplished with passive digital signage (such as plasma screens). However, it's important to note that unlike passive digital signage, touch screens also create information that can be sent back to content creators. The measurement of touch feedback patterns can give content creators immediate access to the reactions of their customers, which in turn can be used to enhance the content. For instance, one can imagine a marketing firm creating three different wall applications to advertise a new movie. These three applications can be distributed to three different test markets. After some time period, the marketing firm can use touch information to gauge the effectiveness of each application and distribute only the best application to all markets.

3 Challenges and Approaches

Having laid out some of the potential of the interactive wall as a medium, we will present some of the challenges which arise in developing content that uses the medium effectively. As we do, we will present various approaches to working through these challenges.

3.1 Touch Interaction and Multiple Users

The fact that a wall-size display can be interactive raises a number of questions for designers: how will people know how to interact with it? What will their initial expectations be, and will these be met? What sorts of places are suitable for an interactive wall?

Touch interaction going beyond simple button-pressing is somewhat novel, but it can also be quite intuitive for most users (once they recognize that the wall is actually touchable). In our system, we chose to recognize only the simplest touch inputs: touching, releasing, and dragging objects on screen with a finger or a cluster of fingers. We also support "waving away" certain items as our only supported "gesture", but the gesture is optional. Requiring users—particularly casual users in a mall, airport, street environment, or museum—to learn a vocabulary of special gestures in order to interact sets the barrier to entry needlessly high. Specialized gesture vocabularies might be suitable to a command and control environment (though we question this as well); in an entertainment context, they are probably out of the question.

The potential for simultaneous user interactions presents a new dimension for user experience—and a new design challenge. The same questions as above must be asked again in a different light: will people realize that the wall can handle multiple touches? Operating on the assumption that they are dealing with a computer—and computers have a single input point—users might assume the contrary. (A user study on questions like these would be very illuminating.) What kinds of applications will engage multiple people, and not just a single user?

There are very few standing guidelines for multi-user interactive experience, except, again, in the area of arcade games. From arcades, we do know that engaging applications, especially those that support and reward constant communication between players, can succeed. It is also important to keep learning time to a minimum, both so that newcomers feel entitled to try the game, and so that the onlookers—who themselves are crucial to attracting interest, and attracting new players—stay engaged, instead of dissipating in boredom as a new player tries to gain enough mastery to actually play the game. When no one is playing, arcade games typically feature an "attraction loop", which demonstrates various important facts: that the game is available to be played; what kinds of experiences are available in the game; and even how to play. We may want to offer a similar functionality for wall applications.

From the software design perspective, we have added a layer of complexity by accommodating multiple users, each of whom might be performing multiple touches, all within the scope of one application. Most UI tools and packages are tuned for a situation in which one event arrives at a time, in a coherent sequence that representing the behavior of a single user. We are unaware of any standard applications in which two or three mice simultaneously operate in a single application space, but that is precisely the situation that the multi-user interactive wall creates for the designer. Compounding this is the fact that unlike mouse pointers, which have an identity and a given position at all times, touches are ephemeral, disappearing and reappearing; it now becomes a task in itself to stitch together various touch signals into a coherent picture of user behavior.

These considerations rule out certain styles of interface behavior, such as, for example, anything that is modal—i.e. any interface element which assumes it can get, or has got, exclusive control over the environment. Menus, for example, are modal in most interfaces: a click opens the menu; another click, on a menu item, selects that item. What about clicks elsewhere on screen? They close the menu—in fact, that is the principal means of dismissing a menu we don't wish to use. In a multi-touch environment, this logic (now a part of every GUI user's muscle memory) is probably unsuitable, because events that are distant from me on screen should not be able to affect the state of "my" interface widgets.

In fact, most of the classic interface widgets that are the standard components of a modern desktop GUI are built on a similar logic. If we wish to use them, they need to be rewritten so that they do not assume a single-user environment. But even if this is done, some of these widgets are unsuitable for the wall interface in the first place. Some of the reasons for this will become clear in our discussion of the challenges of the wall's large size.

The assumption of modality, however, is not restricted to GUI widgets. It is a core principle of desktop software design, and we must think carefully about how to avoid accidentally re-invoking this principle in a context where it is inappropriate.

In Figure 1 (adapted from content at [6]), we show three possible software environments. The left panel of Figure 1 depicts a number of buttons arranged along the top of a screen. One of the buttons has been selected, causing a product detail window to appear below. The scheme is typical of what we might find in a Macromedia Flash application on a website. Note that the menu system takes up the whole width of the application environment. This reveals one kind of modality assumption: it assumes exclusive rights to the entire upper edge of the (single-use) viewing port. The product detail window takes up the rest of the space, again modally, covering up whatever might have previously occupied that real estate. This kind of logic works well in a single-user environment with a small screen, because it uses the limited screen real estate effectively.

Figure 1. Layouts in different form factors.

In the center panel of Figure 1, the same kind of scheme has been naively applied to an interactive wall. The buttons are still ranged along the top of the screen—but this means that now, each user is closer to some of them, and farther from others. This probably doesn’t make sense if both users might like to work independently. And of course, the fact that the product detail window occupies most of the wall surface is an even greater impediment to multiple use.

In the right panel of Figure 1, the menu system has been removed—or rather, the menuing function has been spread across the canvas, in the form of a product matrix. The users can now browse independently, even though there is a single product matrix which they must share. They simply browse different parts of the matrix, without preempting one another visually. In this particular application, each user does not have total access to all the products from where he is standing; each product only exists at one spot on the matrix. But that may be a reasonable design choice in this case; perhaps as a user physically moves rightwards, along the wall, they move towards higher-priced cars. The structure of the wall space reflects the continuum of products. Now, moving one’s body is a meaningful action within that space, rather than an arbitrary artifact of a UI widget. Software design in these kinds of scenarios might actually incorporate a notion of crowd flow.

3.2 Designing at Scale

Building compelling applications for a canvas the size of a wall presents a challenge which may surprise many designers. Over the past few decades, just as most interface logic has been optimized for a single pointer, as mentioned above, most GUI design has assumed a relatively small, well-bounded interaction surface. When the surface is small, we can assume that the whole space randomly accessible to the user's eye, at almost no time cost. This lets us put a clock in the corner, for example, for handy reference. What's less handy, however, is when the clock is eight feet away, and I have to move my whole body in order to read it. As the cost of viewing access goes up, a static viewing target like the clock loses some utility.

Similarly, in a compact desktop environment, designers take advantage of the small screen size and the user's muscle memory when they create static click targets, such as menu bars, that the user can visit with their mouse pointer again and again, from anywhere onscreen, at a low navigation cost. They place these targets at the edge of the screen when they can, to help the mouse pointer hit the target without overshooting. But in the new medium of an interactive wall, most of these optimizations must be undone, since the "cost structure" they are based upon has to be refigured.

For one thing, the edge of the screen is not privileged as a click target any more—the opposite is true, since the center of the screen is usually much more accessible from the average hand position. The existence of multiple users might mean that we need multiple copies of a common click target. Or, alternatively, we might do away with static click targets entirely, replacing them with controls that follow the user somehow, either by appearing in appropriate contexts, or by hovering near the area of the user's interactions.

3.3 Fleshing Out the Field of View

Almost by definition, a media wall that is touch-interactive must have a high resolution—that is, a high pixel density—because it will be experienced from arm's length, and we want to ensure significant quality and detail at that close range. Surprisingly, however, this is not always the case. We believe that designers should present the user with a rich visual experience *as they use the wall*, not afterwards, when they step back. Of course, many command center installations, projected presentations, and other large displays are only intended to be viewed from across a room, and consequently feature a fairly low pixel density. But at arm's length, those same densities translate into a sharp limit on the quality of the user's visual experience.

In our experience, single-projector interactive systems (such as Smartboard and similar solutions) suffer from a field-of-view problem—they expect users to touch the screen, but they do not offer a useful amount of visual information at that range. (See the left-hand side of Figure 2.)

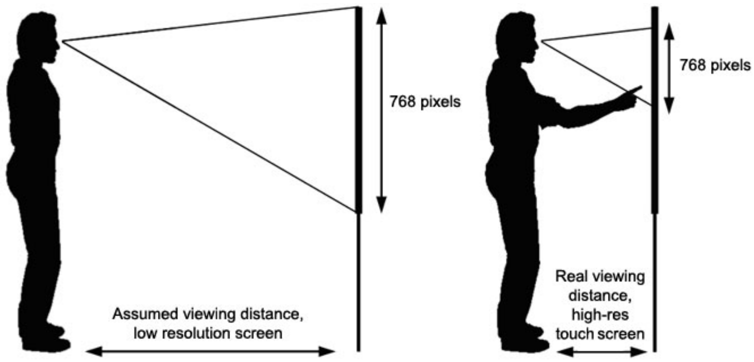

Figure 2. Necessary visible pixels as a function of distance.

We want to include enough visual information, and enough functionality, in a human see-and-touch range (perhaps three or four feet across), that a person can do something interesting, without having to walk around too much, or step back from the wall; all this requires high pixel density. Thus, the way we author any one area of a wall application, and the interactive elements we embed in the design, will often be constrained by the human viewing angle (as well as others physical factors, such as height and arm-span). Density is also the feature that allows multiple users, interacting simultaneously, to have access to a reasonable amount of content even as they stand somewhat close together. As a rule of thumb, we feel that there should be a laptop screen's worth of pixels (1024 by 768 or so) in the user's vertical field of view as they stand at interaction distance, as shown in the right hand side of Figure 1.

But at the same time, we must expect users at all distances, at all times; the same content which might seem suitable for presentation at a distance must also stand up to the scrutiny of users who are much closer in. Imagine a wall application for a car show. The designer might put a large interactive video in the center of the wall, and smaller stations for interactive product information on either side. The interactive video is operated by a single user, but it is wide enough (perhaps six feet wide) that it can be seen by passersby.

If we were to repurpose standard DVD-quality video for this application, and simply made it "big", it would probably not be effective. Passersby would be satisfied, because of the small portion of their field of view occupied by the (relatively) small amount of visual content. But for the primary user—the one who is running the show and creating the spectacle—there isn't much visual pleasure in the experience. The effect is similar to standing too close to a TV. And for anyone interested enough to approach the wall, visual quality falls off quickly. And, in the case where the wall is not in use, running an attraction loop, the lack of visual quality may discourage interaction in the first place. If it looks like a big-screen TV, people will probably assume that it is one.

Broadly speaking, there are many situations in which the proportions of ordinary digital graphics should be reconsidered for wall-based designs. Figure 1 showed a web application with a matrix of cars [6]. The web application is optimized for an 800x600 window. To fit in the available space, each car is represented by the low fidelity graphic shown on the left side of Figure 3. On a 2048x1536 wall display, the same matrix can be represented with a much larger number of pixels. If one keeps the relative sizes of all of the elements the same, each car can be—should be—represented a higher-fidelity version such as the one on the right hand side of Figure 3.

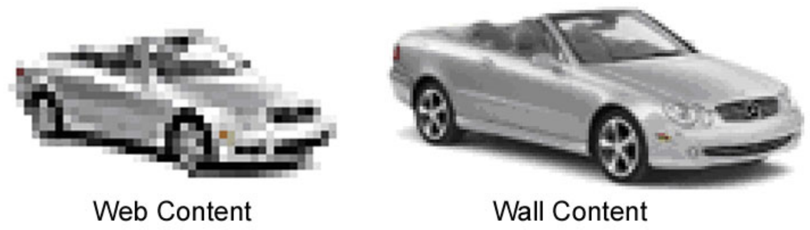

Figure 3. Different amounts of detail at different pixel densities.

The possibility for much higher fidelity—which is also a demand for that fidelity—can create interesting challenges for content acquisition and creation, as all creative assets have to be created at different scales and resolutions than do web or print graphics. However, as Figure 3 shows, better art assets create better product depictions and rich user experiences.

4 Conclusion and Examples

In the previous sections, we have described some of the advantages of large format displays as they pertain to entertainment and advertising, we have discussed some of the challenges, and we've used abstract examples to describe some of our current thinking about how to approach these challenges. We are still very early in the process of understanding the medium and refining our design choices, and we hope that these ideas will help others think about how to best design content for large screens.

As with early word processing and web design, the first examples of content designed for a new medium are not necessarily the best indicators of the eventual potential of that medium. Having said that, Figure 4 shows one example of what we have designed for our 4096x1536 touchable display. The depicted application is aimed at multiple users in a public environment, provides a rich visual experience at various distances, and is driven by intuitive touch interaction.

We are using this application, and others like it, to explore the medium and define approaches that work best for the form factor. These applications also serve as a basis for user studies into the usability of large, public, multi-user displays. We hope to present the results of those investigations in the future.

Figure 4. One example of an interactive advertising and entertainment application.

References

1. Jupiter Systems. http://www.jupiter.com
2. G. Humphreys, M. Houston, Y.-R. Ng, R. Frank, S. Ahern, P. Kirchner, and J. T. Klosowski, "Chromium: A Stream Processing Framework for Interactive Graphics on Clusters," presented at SIGGRAPH, San Antonio, Texas, 2002.
3. Li, K., Chen, H. et al. Building and Using A Scalable Display Wall System. IEEE Computer Graphics and Applications, Volume 20 , Issue 4 (July 2000) Pages: 29 - 37
4. Dempski, K., Harvey, B.: Supporting Collaborative Touch Interaction with High Resolution Wall Displays. 2nd Workshop on Multi-User and Ubiquitous User Interfaces (MU3I) January 9, 2005 at IUI 2005
5. NCSA Tiled Display Wall. http://brighton.ncsa.uiuc.edu/~prajlich/wall/
6. Mercedes-Benz. http://www.mbusa.com/brand/selector/controller.jsp?filter=msrp