A version of this article will appear in
IEEE Internet Computing, November/December 2001
Special Issue on Personalization and Privacy

Personalization and Collectivization at the Intersection of the Virtual and Physical Worlds

Joseph F. McCarthy
Accenture Technology Labs
3773 Willow Road
Northbrook, IL 60062 USA
mccarthy@cstar.accenture.com
http://www.accenture.com/cstar/

Abstract

The personalization of an individual computer user's experience of the virtual world, especially with respect to one's interactions with the Internet and World Wide Web, has received a great deal of attention. However, many of us continue to spend most of our time in the physical world, and much of this time is spent in the company of others. We are exploring how we might best take advantage of what the virtual world has to offer in specific contexts in the physical world. In particular, we are developing systems that adapt physical spaces by allocating influence over certain aural and visual aspects of these spaces collectively to the group of people inhabiting those spaces at any given time.

Introduction

One of the most exciting new capabilities engendered by the World Wide Web is *personalization* – the ability to customize each individual user's experience of electronic content. This customization can take place along a number of dimensions:

- different sources of content (e.g., different news organizations or topics)
- arrangement of content on a screen
- delivery mechanisms – system-initiated ("push") or user-initiated ("pull")
- delivery vehicles (web browser, mobile phone, pager, FAX printer)

Never before have we had a medium that was both so widely available and so individually customizable.

A number of systems have been developed to exploit this capability for personalizing the web; for example, BotSpot (bots.internet.com) lists over 60 personalized agents that are designed to aggregate online news stories based on a user's interest profile (and 15 more that focus exclusively on Usenet newsgroup articles). In addition to "the usual suspects" – news, weather, stocks and sports scores – a number of other types of content can be personalized, such as recommendations for books (www.amazon.com) or movies (movielens.umn.edu).

While the efforts to personalize web content are having a significant impact on how people seek – or, in the case of "push" technology, are sought out by – information in the *virtual* world[1], we believe that we are on the brink of a new generation of personalization, in which aspects of the *physical* world can be customized. We are entering an era of ubiquitous computing [Weiser & Brown, 1997], in which we will have computing and communication capabilities available in all kinds of environments and situations beyond the "traditional" model of sitting at a desktop computer workstation (and beyond even the more recent interaction models incorporating handheld devices and mobile phones). Under this new model of computing, human-computer interaction issues start to evolve into issues of *inhabitant-environment interaction*.

Most ubiquitous computing applications focus on how these new capabilities will affect an *individual*'s interactions with his or her environment. However, much of our time is spent in shared physical spaces, so it is important to consider how an environment might effectively sense and respond to *groups* of co-located people. This paper raises a number of issues for what might be called "UbiGroup" applications, and describes research that addresses some of these issues.[2]

Many of the computer applications we are familiar with today are used in the *foreground* of our attention, e.g., purchasing a book on your desktop computer's browser or checking stock quotes on your web-enabled telephone. However, as computers become smaller and cheaper, they will increasingly be embedded in a variety of objects that do not typically require, and often do not permit, our focused attention. Such applications will operate in the *background*, at the periphery of our attention; they will affect aspects of our environments, but not necessarily assist us directly with the task(s) at hand.

One example of this type of application is MUSICFX [McCarthy & Anagnost, 1998], a system situated in a fitness center environment, that is aware of who is working out and what they generally like to listen to, and uses this knowledge to determine the best music to play at any given time. Another example, which moves from the workout place to the workplace, and from the aural domain to the visual domain, is GROUPCAST [McCarthy, *et al.*, 2001], a system situated in a common area of the workplace environment, that is aware of who is walking by, what their interests are, and displays content on a large monitor that is of mutual interest

The remainder of this article will present these two applications in greater detail, highlight some of the personalization (and privacy) issues that arise in the customization of physical spaces, and conclude with a discussion of some future applications and extensions.

---

[1] We will refer to content that is available primarily in the electronic realm of the Internet and World Wide Web as the "virtual world," as contrasted to the tangibility of the physical world; some people refer to these as the worlds of "bits" and "atoms." We also wish to contrast "virtual world" with "virtual reality" which is a specific type of virtual world which seeks to recreate as many aspects of the physical world as possible.

[2] We note that there are a number of systems that take into account groups of people who are not collocated, such as the Internet radio station www.launch.com, which seeks to create virtual communities of listeners on the web, and a new feature added to MovieLens (movielens.umn.edu), where people can search for movies that a group of people might enjoy together [O'Connor, *et al.*, 2001] – though these people need not be collocated when the search is undertaken.

MUSICFX

Any intelligent system that adapts to the preferences of its inhabitants needs three main components: a mechanism for detecting inhabitants and their activities, a representation of inhabitant preferences, and an algorithm for deciding how to adapt based on those preferences.

MUSICFX detects inhabitants by requiring members to login, using a proximity badge reader connected to a desktop computer and standard-issue employee badges, as they enter the fitness center. Rather than requiring a member to explicitly logout – for which there exists no significant incentive – we set an expiration timeout of 90 minutes after each login, after which time the system presumes the member has left.[3] MUSICFX assumes its inhabitants' activities can be broadly classified as exercising, and doesn't require finer-granularity distinctions.

The MUSICFX preference database represents members' ratings of each of 91 genres of music, each available on a separate station from a satellite music service. Each genre is rated on a 5-point scale, from +2 through –2, interpreted as "I {love, like, don't care, dislike or hate} this kind of music." The initial set of preferences is submitted to the system remotely via an electronic enrollment form; members can update these preferences in the fitness center at any time, using an interface on the computer to which the badge reader is attached.

When a member logs in to the system, that person's preferences are retrieved and added to the current pool of preferences. The MUSICFX Group Preference Arbitrator sorts the list of genres from most popular to least popular, and then uses a weighted random selection algorithm[4] to select one of the most popular genres to play. The Arbitrator is invoked each time a person enters or leaves the fitness center, each time a person updates his or her preferences, each time a fitness center staff member adjusts a system parameter, or after a maximum play time for a single genre has been exceeded.

MUSICFX has been running continuously in the fitness center at Accenture Technology Park in Northbrook, Illinois, since November 1997, with over 800 members now enrolled in the system. The system has been very popular, with over 70% of members reporting they like the MUSICFX-controlled music selection better than the previous human-controlled music selection. The benefit most often cited by fitness center inhabitants is the satisfaction they feel in being able to *influence* an environmental factor that affects them all – without having to complain or otherwise attempt to exert more direct *control* over that factor.

---

[3] For convenience, we will refer to the virtual "logout" events that occur 90 minutes after entering a fitness center as members "leaving" or "exiting" the center.

[4] Rather than always selecting the most popular station, which could result in a tyranny of the majority, an element of randomness – with probabilities distributed according to popularity – was introduced in order to inject a degree of variety (& equitability). This way, when the same people work out at the same times, they still get some variety, and people whose preferences lie outside the mainstream, may still hear music they enjoy on occasion.

Personalization & Privacy Issues with MUSICFX

Most of the personalization and privacy issues concerning MUSICFX revolve around the musical preferences of fitness center members. As mentioned above, members submit their initial set of preferences via a Lotus Notes form (Notes is our firm-wide electronic mail system). These forms contain a text field for the fitness center Member ID, a matrix of 91 (stations) x 5 (preference ratings) radio buttons and a text field for comments. Submitted forms are stored in a database, from which they are imported into the MUSICFX preference database (a Microsoft Access database stored locally on the computer in the fitness center). Although the fitness center staff – as well as the system developers – have access to these preferences, there is typically no need for anyone to ever view these forms.

Interestingly, a Perl script that was used to collect the contents of all the “Comments” fields from these forms revealed that several members thought it important to assure any potential viewers that their preference forms reflect their musical tastes *only while working out* (e.g., “I really do like opera, just not while I’m on a treadmill”). This raises the more general issue of *context* for personalization systems, for both the virtual and physical worlds. An online book recommendation system may experience difficulty if a user sometimes orders books for other people, e.g., a partner/spouse, child, parent or friend. As the capability for influencing physical spaces becomes more generally available, one can envision the complexity in specifying contexts for which different preferences apply, e.g., music for working out, working, eating, drinking or romancing (not to mention contexts in which several activities may be taking place simultaneously).

Within the fitness center context itself, a number of privacy issues arise. Although the forms are submitted in a manner that shields the preferences from other members of the fitness center, when people update their preferences on the computer in the fitness center, other members looking over their shoulder may be able to see some of their preferences. This has generally not been a problem, since people tend to update their preferences in a fit of anger (a member saying “I can’t stand this music!” on the way from the elliptical trainer to the computer to change his or her preference to “hate” for the currently playing genre). It should be noted that this over-the-shoulder disclosure may also exist for personalization systems that operate exclusively in the virtual domain, depending on where people access content (e.g., displaying book recommendations in a shared space with several desktop computers adjacent to one another).

Other privacy issues arise in the context of a system that controls a resource shared – and influenced – by a group of collocated people. As soon as a new member enters the fitness center and badges in to the system, if the currently playing station falls out of favor (say, because this person hates that station, and it was not among the most popular to begin with), the system changes the station – to one that is likely to be preferred by that person (as well as the rest of the group). Initially, we considered introducing a random delay – in the range of 1 to 2 minutes – between the time a person badges in and the time the system changes the station, introducing a “plausible deniability” to station changes so that we could shield inhabitants from revealing anything about their musical preferences. However, we decided that individual fitness center members might like to know if/when the station changes as a result of their arrival, since if a station they do not like [any longer] is selected, they might be immediately "notified" that their preferences need updating and can take action on the spot. This scenario has been observed by the fitness center staff on a few occasions, where someone badged in, an unusual station was selected, and the person went over to the MUSICFX computer with a confused expression and changed his or her preference rating, resulting in a new station being selected. Of course, the drawback is that people may feel pressured into changing their preferences to accommodate [their perception of] more common or popular tastes, if the music suddenly changes to, say, Hawaiian Melodies – which may be popular in some contexts, but probably not in a fitness center. In practice, since the system has a bias toward stations that are popular with the current group of people, this has not been a problem.

GROUPCAST

After our success in creating an adaptive environment in the workout place, we shifted our attention to the workplace. A physical space that can sense people in the vicinity, and has knowledge of their interests, can use this information to create new informal interaction opportunities for these people. For example, a shared public display in a workplace, combined with a tracking system, can display information of mutual interest to the people passing by the display. People may choose to take advantage of this information to initiate a conversation with someone about whom they may know very little, leading to an increased sense of community in the workplace [Deutsch, 1995; Naylor, et al., 1996; Putnam, 2000].

Other researchers have investigated how to create greater awareness among people who are electronically connected but not physically collocated [Erickson & Kellogg, 2000; Zhao & Stasko, 2000; Sawhney, *et al*., 2001; Greenberg & Rounding; 2001]. Our focus has been how to create greater awareness of each other when people are gathered together – or passing each other – in the same physical space. We also want to distinguish our work from other work using large public displays in the foreground to support the performance of primary work activities (e.g., Streitz, *et al*., [1999]); although GROUPCAST uses a large public display, it is intended to be more of a background or peripheral display, and we believe that the content is more likely to spark informal conversations if it is not directly related to work activities.

As an example scenario of GROUPCAST in action, suppose Joe and Teresa pass each other in the hallway fairly regularly and yet know very little about each other. However, the "Wine of the Day" web site pops up as they both pass by a GROUPCAST display, leading to a spontaneous and serendipitous discussion about the merits of old-vine zinfandels (see Figure 1). After the discussion ends, they both go away, knowing a little more about each other, and, assuming the discussion did not deteriorate into a vicious argument, they are more likely to have conversations (on wine and other topics) in the future.

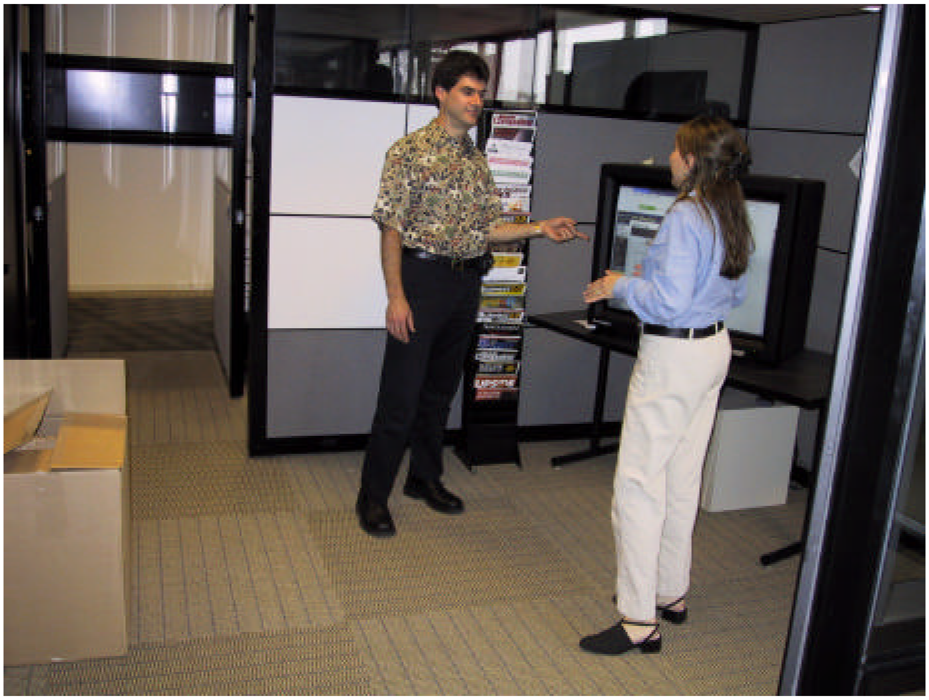

Figure 1: GROUPCAST in context of use.

One of the stumbling blocks we encountered in the initial design of GROUPCAST was how to acquire content that would be of mutual interest to people. We first considered using a large web-based form displaying content areas that people could rate with respect to their interest level (this was the approach we took with MUSICFX, which adapted aural aspects of the workout place rather than visual aspects of the workplace). However, we soon discovered we had conflicting goals: having a profile that would be broad enough to include content of potential interest to a large number of people, and yet still be small enough so that we could reasonably expect people to specify that content, e.g., by filling out a form. By the time we had amassed enough potential content in our profile form, we were fairly confident that no one (besides those working on the project) would take the time to fill it out.

Fortunately, at the same time we were struggling with this issue in GROUPCAST, another application was under development – UNICAST – that would display content on a peripheral display inside a person's office. After we launched UNICAST, we had an insight: if we were to use the UNICAST profile for GROUPCAST, we can rely on people's own self-interest in customizing content that they will see regularly (in their office), rather than struggling with the somewhat less rewarding task of specifying content that only is available when they are in a public area. Since UNICAST content figures prominently in GROUPCAST, we will describe the application, highlighting its context of use and profile information, in the following section.

UniCast

UNICAST is an application that allows users to specify content they would like to see on peripheral displays located within their primary workspaces. In some respects, UNICAST represents an extension of the functionality provided by the PointCast system [cf. Franklin & Zdonik, 1998], which allows people to specify news topics and stock symbols about which they would like to stay informed while their desktop computer is in screensaver mode. However, UNICAST is different in several key aspects: it runs continuously on a dedicated, peripheral display; it allows for a broader selection of content; it reacts to the location of its “owner” via an infrared badge system; and it is tied into and makes use of content belonging to other UNICAST user profiles.

Our model of interaction is primarily that of an ambient display [Weiser & Brown, 1997; Wisnewski, *et al*., 1998; Redström, *et al*., 2000] rather than the primary workstation display used for supporting a user’s primary work tasks – see Figure 2. The hypothesis is that UNICAST content should be interesting, but not terribly important or urgent, since important or urgent information is (or could be) sought out directly on the primary workstation. For example, the author uses UNICAST to cycle through his favorite on-line comics (among other types of content), which help to brighten his day, but are rather peripheral to his work, and which he therefore rarely seeks out on his primary workstation.

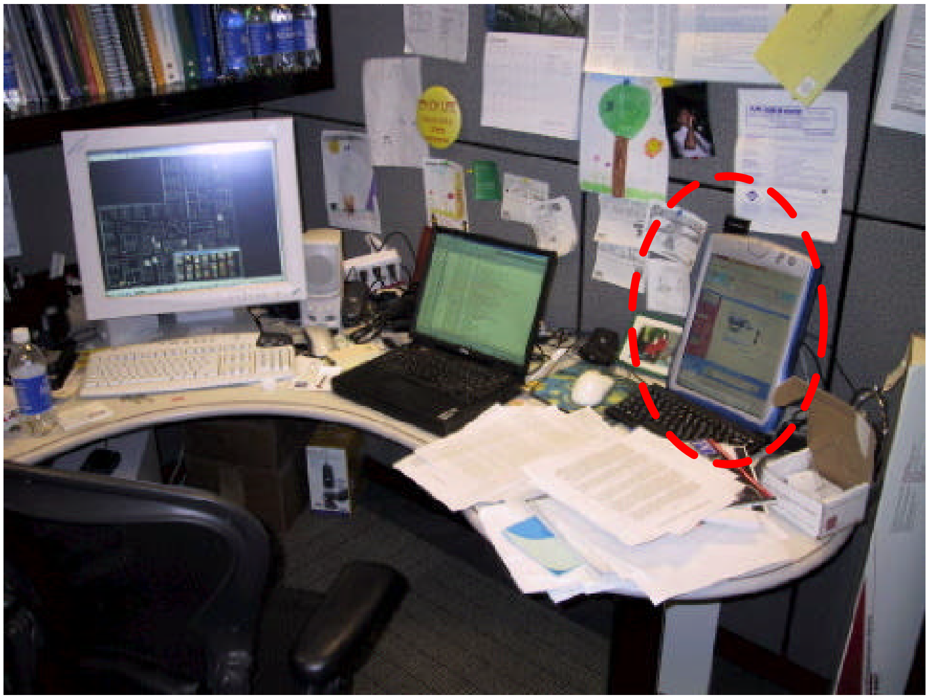

Figure 2: UNICAST display within one office context.

The content for UNICAST includes some of the usual suspects – headlines, stock and weather information – as well as many other types of content. The current implementation includes user-configurable modules of fifteen different classes:

- *Headlines*: The top 5 headlines from any of 273 channels in 16 categories from Moreover (http://www.moreover.com/).
- *Stocks*: Any stock symbol available from Yahoo! Finance (http://finance.yahoo.com/). At most 5 stocks can be displayed per page; for modules with more than this limit, UNICAST randomly selects 5.
- *Weather*: Weather information for any US zip code available from Earthlink's weather portal (http://www.earthlink.net/).
- *Traffic*: Chicagoland Expressway Congestion map provided by the Departments of Transportation of Illinois, Indiana and Wisconsin – IDOT, INDOT, and WISDOT – in conjunction with the AI Lab at the University of Illinois, Chicago's Electrical Engineering & Computer Science Department (http://www.ai.eecs.uic.edu/GCM/CongestionMap.html)
- *Horoscopes*: Any of the 12 signs of the zodiac available from Yahoo! Astrology (http://astrology.yahoo.com/).
- *Web pages*: Any URL specified by the UNICAST user.
- *InfoShare*: Any URL specified by *any* UNICAST user; these URLs are organized by topic and/or people – either the person who posted the content or the person(s) for whom the content was posted (e.g., "I think Tony would be interested in this page") – allowing access by any of three "channels" (this is essentially used as a local "What's Cool?!" repository).
- *Announcements*: Text messages submitted through a special interface that includes a title, body and expiration date; this is mostly use for conference and workshop announcements. Subscription is "all or nothing."
- *Reminders*: Visual and aural reminders of regularly scheduled events for the entire group. Events are individually selectable.
- *WebCams*: Live content from any of 11 Axis 2100 Network Cameras positioned in public spaces throughout the CSTaR area (the video streams from these cameras can only be viewed behind the Accenture firewall, and the URLs are not circulated outside of our group).
- *In/Out List*: A list of who is in the office today (as well as when and where they were last seen) and who is away, based on information gathered through our infrared badge system. People to track are individually selectable. This module is labeled "ActiveMap" for historical reasons (see McCarthy & Meidel [1999] for more information).
- *Factoids*: 363 different "factoids" manually collected from various sources on the web, organized into 8 categories (Culture, History, Human Body, Nature, Recycling, Science, Space and General). Subscription is by category.
- *Flashcards*: Short questions and answers; a default set of US State Capitols augmented by whatever flashcards people add.
- *Artwork*: 1000 images, organized into 10 categories, from Corel's Super Ten Royalty-Free Art Photo Pack. Subscription is by category.
- *Pictures*: Digital images uploaded to a shared directory. Images are individually selectable.

Personalization & Privacy Issues with GROUPCAST

The current implementation of GROUPCAST utilizes all the items in a person's UNICAST profile, with no mechanism for people to specify whether they want any of this content to be displayed in a public space when they are nearby (possibly with others in close proximity). This is clearly a shortcoming that we plan to rectify in the near future. The lack of privacy protection is likely to promote some self-censoring, as people may be unwilling to add content to their UNICAST profile that they are truly interested in, but about which they may be embarrassed in a public setting. Once we add this feature, we hope to learn more about how much of their interests people are willing to reveal (in this context).

One problem with GROUPCAST is that it looks for content in the intersection of people's profiles, and this intersection is often quite small, resulting in little (or no) variety in the content displayed in a public space. A possible solution is to select content randomly from each individual's profile, cycling through different profiles. It may be that this would add more conversation opportunities than using intersection, since if GROUPCAST shows content that is already shown in your own office (via UNICAST), it isn't as interesting – or as likely to provide conversation-starting opportunities – as new content appearing in a public space that is selected from someone else's profile.

Another solution we are considering would be to add a new feature to GROUPCAST for passively collecting content based on what web sites people are visiting in the browsers on their primary workstations (e.g., using an instrumented browser like Curious Browser [Claypool, *et al.*, 2001], or some kind of intermediary to collect URLs, such as WBI [Barrett, *et al.*, 1997]). If we had access to such content, then GROUPCAST could show web pages that the people near the display had all viewed recently. Of course, depending on the nature of those web pages, this could either lead to interesting conversations or people fleeing the scene (or in extreme cases, the termination of employment). We might want to intersperse randomly selected content, to provide plausible deniability, and/or introduce filters for "appropriate" content – though that may come at the expense of some of the most interesting content. Given the goal of creating conversation opportunities, we might just be able to select content without any regard for who is near the display, e.g., what are people reading/viewing lately?

Finally, we might want to utilize some of the aforementioned "InfoShare" content, whereby one user (poster) specifies content that is likely to be of interest to another user (postee). Currently, this content becomes available only on UNICAST displays in individual offices, but it may be that having this content displayed in a public space when the poster and postee are both nearby, creating a very targeted conversation opportunity ("Oh, I wanted to show you this article…"). This feature might also bring up some privacy concerns: people might not want others to know what they are posting to whom. However, this concern also exists in the virtual domain, since people who send around URLs to others in an email message might have that message forwarded on to others.

Future Adaptive Environments

We have thus far focused primarily upon MUSICFX and GROUPCAST, and some short-term extensions to these applications in their current contexts of use. In this section, we will expand the discussion to envision how other kinds of environments might effectively adapt to their inhabitants, other ways that inhabitant preferences might be collected, and what kinds of personalization and privacy issues might arise in such contexts.

Automatic adaptation of music playing in an environment would be useful in any physical space in which music contributes to the overall experience. One can imagine how restaurants or retail stores might provide a service wherein frequent diners or shoppers would identify themselves – using a plastic card with a magnetic strip, or perhaps a smart card – in order to influence the selection of music being played. All other factors being equal, people will enjoy their time in a dining or retail establishment more if they are listening to music they like rather than music they don't like, leading to greater customer satisfaction and loyalty, an increase in the number of visits and time spent in the restaurant or store, and greater sales.

Although the prospect of customers identifying themselves to restaurants or stores may raise some privacy concerns, people are already doing so whenever they use a credit card for purchases. In fact, most people are generally willing to forego some amount of privacy for a commensurate benefit: many stores (particularly chains) now have loyalty cards that entitle their holders to small discounts, and an increasing number of restaurants are doing the same. These loyalty cards also act as identifiers, so using the cards to provide additional benefits, e.g., adaptive music, may help to overcome the privacy concerns of people and lead to greater utilization of the cards (providing greater benefits to the merchants as well).

Once a store knows something about the identities, and the purchase histories, of its customers, dynamic promotions can be targeted to the current group of inhabitants while they are shopping. For example, large visual displays could show dynamic advertisements, and perhaps offer discounts, for products that are relevant to that group of shoppers currently in the store, rather than the general population. Likewise, a movie theater could tailors its previews to the people actually in the theater based on the movies those people have already seen (rather than a generic selection of current and upcoming releases). Many theater chains have already instituted frequent moviegoers' programs, and adaptive previews might provide additional benefits for joining such a program.

Airports could have displays that adapt to the people who have checked in but not yet boarded their planes; airlines already know something about the destinations their customers have visited – or are planning to visit – and many airlines also offer their own credit cards, which could provide a great deal more information about the interests and preferences of their passengers. One of the more appealing aspects of this scenario is that adaptation would be based on *collective* histories, interests and preferences; public displays that adapt to detailed information about an *individual* would likely be viewed as a much greater privacy intrusion by many people.

Of course, we need not limit ourselves to air travel – other modes of transportation provide additional opportunities for environments that adapt to groups of people. For example, one can imagine billboards along a highway adapting to the drivers passing nearby. This may be particularly useful in areas with toll roads, where electronic toll-paying devices in cars could be subsidized by advertisers who can then tailor their advertisements to nearby drivers, in much the same way that the aforementioned supermarket could tailor in-store advertisements to the current group of shoppers. Similarly, advertising spaces in trains and buses might adapt to their passengers, based on their use of an identification card such as a monthly pass, assuming that other useful information could be collected about such passengers. In such a scenario, government-subsidized public transportation could transition to advertiser-subsidized public transportation.

Conclusion

There are many benefits to be derived from personalization of electronic content in the virtual world; however, we believe there are even greater benefits to be derived from personalization as it is applied to new contexts – beyond the desktop – in the physical world. We have described two research projects in which we have started exploring benefits that can be realized when people are willing to reveal something about their preferences and interests in the workplace and the workout place, and briefly described several possible extensions of these ideas into other contexts.

Utilizing inhabitant preferences and interests in the *physical world*, particularly in public spaces, in some ways intrudes on people's privacy more than it would in the virtual world: it is easier to maintain different addresses, aliases or virtual personae in the virtual world than in the physical world, so people can more effectively shield or mask details about themselves in their online interactions. However, the use of these preferences and interests in a *group context*, wherein there exists some degree of plausible deniability ("That lingerie advertisement is clearly not based on *my* preferences!"), may diminish people's concerns. And, if we can provide sufficient benefits – think of a world without "elevator music" or other lowest common denominator environmental factors – we believe that people will not only accept but embrace the technologies that will make adaptive environments possible.

Acknowledgements

The author gratefully acknowledges the contributions made by colleagues to the applications and the ideas described in this article, including Theodore Anagnost, Tony Costa, Jeremy Goecks, Elaine Huang, Edy Liongosari, Eric Meidel, and Joe Tullio.

References

Barrett, Rob, Paul P. Maglio and Daniel C. Kellem. 1997. How to Personalize the Web. In *Proceedings of the 1997 Conference on Human Factors in Computer Systems (CHI '97)*, Atlanta, pp. 75-82.

Claypool, Mark, Phong Le, Makoto Waseda and David Brown. 2001. Implicit Interest Indicators. In *Proceedings of the International Conference on Intelligent User Interfaces Conference (IUI 2001)*, Santa Fe, New Mexico.

Deutsch, Claudia H. 1995. Commercial Property; Communication in the Workplace; Companies Using Coffee Bars to Get Ideas Brewing. *The New York Times*, 5 November 1995.

Erickson, Thomas and Wendy A. Kellogg. 2000. Social Translucence: An Approach to Designing Systems that Support Social Processes. *ACM Transactions on Computer-Human Interaction*:7(1):59-83.

Franklin, Michael, and Stan Zdonik. 1998. Data in your Face: Push Technology in Perspective. In *Proceedings of the 1998 ACM Conference on Management of Data (SIGMOD '98)*, Seattle, pp. 516-519.

Greenberg, Saul, and Michael Rounding. 2001. The Notification Collage: Posting Information to Public and Personal Displays. In *Proceedings of the 2001 ACM Conference on Human Factors in Computer Systems (CHI 2001)*, Seattle, pp. 514-521.

Herlocker, Jonathan L., Joseph A. Konstan and John Riedl. 2000. Explaining Collaborative Filtering Recommendations. In *Proceedings of the ACM 2000 Conference on Computer Supported Cooperative Work (CSCW 2000)*, Philadelphia, pp. 241-250.

McCarthy, Joseph F., and Theodore D. Anagnost. 1998. MUSICFX: An Arbiter of Group Preferences for Computer Supported Collaborative Workouts. In *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work (CSCW '98)*, Seattle, pp. 363-372.

McCarthy, Joseph F., and Eric S. Meidel. 1999. ACTIVEMAP: A Visualization Tool for Location Awareness to Support Informal Interactions. In Hans W. Gellersen (Ed.*) Handheld and Ubiquitous Computing: Proceedings of the First International Symposium (HUC '99)*, Karlsruhe, Germany, September 1999. Lecture Notes in Computer Science, Vol. 1707, Springer – Verlag, Heidelberg, pp. 158-170.

McCarthy, Joseph F., Tony J. Costa and Edy S. Liongosari. 2001. UNICAST, OUTCAST & GROUPCAST: Three Steps Toward Ubiquitous Peripheral Displays. To appear in *Proceedings of the International Symposium on Ubiquitous Computing (UBICOMP 2001)*, Atlanta.

Naylor, Thomas H., William H. Willimon and Rolf Österberg. 1996. The Search for Community in the Workplace. *Business and Society Review*, 97:42-47.

O'Connor, Mark, Dan Cosley, Joseph A. Konstan and John Riedl. 2001. PolyLens: A Recommender for Groups of Users. To appear in *Proceedings of the Seventh European Conference on Computer Supported Cooperative Work (ECSCW 2001)*, Bonn.

Putnam, Robert. 2000. *Bowling Alone: The Collapse and Revival of American Community*. Simon & Schuster.

Redström, Johan, Peter Ljungstrand and Patricija Jaksetic. 2000. The ChatterBox: Using Text Manipulation in an Entertaining Information Display. In *Proceedings of Graphics Interface 2000*, Montréal.

Sawhney, Nitin, Sean Wheeler and Chris Schmandt. 2001. Aware Community Portals: Shared Information Appliances for Transitional Spaces. In *Journal of Personal and Ubiquitous Computing*, 5(1):66-70.

Streitz, Norbert A., Jörg Geißler, Torsten Holmer, Shin'ichi Konomi, Christian Müller-Tomfelde, Wolfgang Reischl, Petra Rexroth, Peter Seitz and Ralf Steinmetz. 1999. i-LAND: An Interactive Landscape for Creativitiy and Innovation . In *Proceedings of the 1999 ACM Conference on Human Factors in Computing Systems (CHI '99)*, Pittsburgh, PA, pp. 120-127.

Weiser, Mark, and John Seely Brown. 1997. The Coming Age of Calm Technology. In Peter J. Denning & Robert M. Metcalfe (Eds), *Beyond Calculation: The Next Fifty Years of Computing*. Springer – Verlag, pp. 75-85.

Wisneski, Craig, Hiroshi Ishii, Andrew Dahley, Matt Gorbet, Scott Brave, Brygg Ulmer and Paul Yarin. 1998. Ambient Displays: Turning Architectural Space into an Interface between People and Information. In Norbert A. Streitz, Shin'ichi Konomi and Heinz-Jurgen Burkhardt (Eds.) *Cooperative Buildings - Integrating Information, Organization and Architecture: Proceedings of the First International Workshop on Cooperative Buildings (CoBuild '98)*, Darmstadt, Germany. Lecture Notes in Computer Science, Vol. 1370. Springer - Verlag, Heidelberg, pp. 22-32.

Zhao, Qiang Alex, and John T. Stasko. 2000. What's Happening? The Community Awareness Application. In *2000 ACM Conference on Human Factors in Computer Systems (CHI 2000) Extended Abstracts*, The Hague, pp. 253-254.